IBA TECHNICAL REVIEW

8

Digital Video Processing — DICE

IBA
INDEPENDENT
BROADCASTING
AUTHORITY

IBA TECHNICAL REVIEW September 1976
ISSN 0308-423X

8 Digital Video Processing—DICE

CONTENTS

Technical Editor: Pat Hawker, IBA Engineering Information Service

Additional Copies
Subject to availability, further quantities of this IBA Technical Review may be obtained on application to Engineering Information Service, IBA, Crawley Court, WINCHESTER, Hampshire SO21 2QA.
No charge will be made for small quantities.

IBA
INDEPENDENT BROADCASTING AUTHORITY
70 Brompton Road, London SW3 1EY *Tel: 01-584 7011 Telex: 24345*

Introduction

by J B Sewter
Chief Engineer (Development & Information)
Independent Broadcasting Authority

Progress within the field of television technology has become so rapid that any development, heralded as a revolutionary breakthrough in one year, can by the next become standard practice, and within two further years be relegated to obsolescence. Yet, of the new concepts which emerged and proved themselves worthy of adoption in 1972, there is one which has probably outshone all others; and now, at June 1976, it still survives and still remains a clear leader of the technology. This is DICE – the *D*igital *I*ntercontinental *C*onversion *E*quipment designed and developed by the IBA. It has been in full operational use at ITN as a one-way conversion system since Spring 1973 and as a two-way system since Spring 1975.

DICE, in its several versions, has been proclaimed the most advanced single piece of digital equipment either described or shown at IBC 1974, Montreux 1975 and NAB 1976; and, in its latest form, it is ready to make its mark at IBC 1976.

There are, of course, other operational digital-video equipments. Digital time-base correctors and digital frame synchronisers are already making significant impact on American television by opening the way to electronic news-gathering and the easy and convenient handling of remote sources. Each of these important developments uses techniques first shown to be practicable in the early work in the United Kingdom on line-store and field-store standards conversion.

The ability to remain throughout several years the acknowledged leader in any branch of modern technology depends on more than merely fortuitous circumstances. Quite deliberately, as soon as the first part of the conversion to colour broadcasting on uhf was completed, the IBA Experimental and Development Department was set to undertake a long-term study of the likely impact of digital techniques on television. The study included methods for the operation and control of large unattended television networks (from which has emerged DAME – *D*igital *A*utomatic *M*easuring *E*quipment), and for data communications, initially for housekeeping purposes but which soon led to teletext and ORACLE – (*O*ptional *R*eception of *A*nnouncements by *C*oded *L*ine *E*lectronics).

In 1968-69, when those studies began, the likely cost of digital memory appeared quite prohibitive; and it was only by looking ahead and boldly budgeting on extrapolations of cost trends – always a risky procedure, and especially so in time of monetary inflation – that we were encouraged to tackle the development of such systems.

This volume of the *IBA Technical Review* is devoted to a detailed account of recent developments in DICE, including the two-way version with improved interpolation which is now being manufactured commercially under licence by Marconi Communication Systems Ltd.

All this, we believe, is still only a beginning. Digital techniques of the future will probably pervade throughout the studios as well as throughout the transmission links. The whole weight of economic pressures – coupled with the pride taken by engineers in the improved results which become possible with the more rugged digital signals – will ensure that victory will be with the digits. But, engineers are more than crystal-ball gazers. Rather than presume to make heady forecasts of exciting things to come, we prefer to describe factually certain things which have been done.

Already, DICE has established itself in the history of television broadcasting. Its continuing progress provides a significant success story. In the following pages IBA engineers concerned with the project relate that story in engineering terms.

JOHN L E BALDWIN, BSc, Member of the Institute of Physics and of the Royal Television Society, Freeman of the City of London, has been working in television since he joined Rank Cintel in 1950. After 14 years he moved to Peto Scott Ltd (part of Philips Gloeilampenfabrieken) becoming Chief Engineer. In 1967 he joined the Independent Television Authority, which became the IBA on July 12, 1972, as Head of the Video and Colour Section of the Experimental and Development Department. In 1975 he was awarded the David Sarnoff Gold Medal of the Society of Motion Picture and Television Engineers for his personal contributions and as the leader of the team that developed DICE.

Digital Standards Conversion

by J L E Baldwin

Synopsis

The first part of this section describes the background to the development of digital standards conversion from some early thoughts in 1959 to the active development of a line-rate converter in the IBA laboratories during 1970-71 and the realisation that digital video processing opens the way to high-grade, drift-free intercontinental standards conversion.

The second part outlines the principles used in DICE as an introduction to the more detailed explanation of specific arrangements described later in this volume, providing a technical linking and introduction to the complete processes used in both directions of conversion, and the reasons for choosing five-line interpolation.

A standards converter is a very specialised item of equipment but many of the techniques used within DICE – Digital Intercontinental Conversion Equipment – may come to have applications over a much wider field when digital processing of video signals becomes commonplace.

Some of the novel techniques developed for standards conversion have already been adopted in synchronisers and timebase correctors but others still remain to be used outside of converters. For example, DICE includes a digital colour coder and decoder for 525-line NTSC; and a digital colour bar generator; it also effectively includes both two-input and five-input 'knob-a-channel' mixers (these are used in the converter to provide movement and line interpolation respectively). An understanding of these digital techniques is thus of interest to engineers working generally in television.

History

My interest in electronic standards converters was aroused in 1959 when it became clear to me that magnetic tape could be used as a temporary store to duplicate or omit certain lines of a field and to alter appropriately the line duration (British Patent Number 898,723, June 18, 1959), and to duplicate or omit a certain proportion of fields, the duration of each field being altered appropriately. These ideas however could not be achieved in *real time*. About 5 hours would have been taken to convert one hour of programmes from 625 to 405 lines and about $2\frac{1}{2}$ hours from 525 to 625. This disadvantage, which arose from mechanical limitations, was so great that

these inventions remained for me just ideas until some time after I joined the ITA in 1967.

But towards the end of 1969 the price of relatively fast digital storage dropped significantly and it seemed likely that the use of digital techniques in line-standards converters was about to become financially, as well as theoretically, attractive. Some preliminary work was started in the IBA, in the middle of 1970, on two key areas of conversion; namely a line-store and a multiplier. Results were promising; development of digital line standards converters went ahead and an experimental equipment was demonstrated in London to the EBU Technical Committee in March 1971: this converter used four-line interpolation.

We needed to be sure that digital video processing could be used, without major problems, in conditions of high rf field strengths so, in the summer of 1971, the converter was taken to the ITA's VHF transmitting station at Croydon, designed to serve over 13 million people in the London area. No problems arose and the converter was used operationally for a period in place of the usual analogue line-store converter.

The performance was so good, substantiating in practice all the theoretical forecasts, that it became obvious that the next step should be to develop a digital *field* standards converter. The Authority authorised the E&D Department to construct an experimental field standards converter provided that this was as simple as possible, yet capable of proving ideas. It seemed sensible to accept that such a design would convert only in the 525-to-625-line direction and that certain operational limitations would be tolerated (for example that the input signal would always contain a colour burst, a limitation which prevented true monochrome conversions).

In November 1971, development of such a standards converter was started in the London laboratories of what was soon to become the Independent Broadcasting Authority. One year later this experimental model was operational, feeding the whole of Europe with pictures of the November 1972 American presidential election received via satellite at the Goonhilly ground station. By now the converter had been named DICE.

Figure 1 shows this first DICE as originally installed at the London studios of Independent Television News in March 1973. It should be stressed that DICE was used not only for the routine conversion of incoming 525-line news material, usually via satellite, but also for the conversion of video tape made in the 525-line standard. Such material included many spectacular entertainment programmes, where good technical quality of conversion is vital, much more so than for news material.

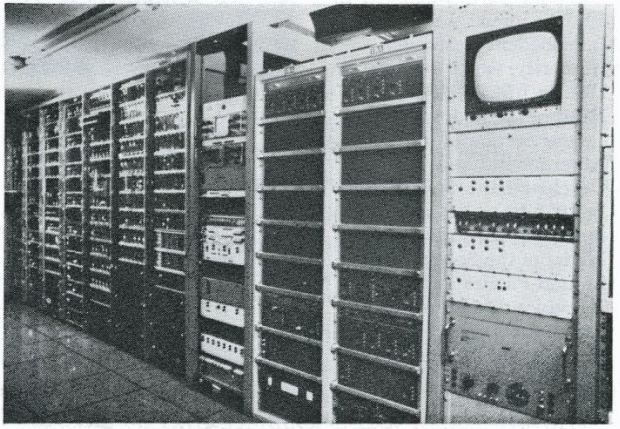

Fig. 1. The prototype one-way DICE was installed in the London studio centre of Independent Television News in March 1973. In this illustration the two racks of DICE (towards right-hand edge of picture) are seen installed alongside the seven racks of an earlier analogue-type electronic standards converter.

The experimental model was capable of operating only in the 525-to-625 direction and was originally intended as a feasibility study and not as a high performance machine. Nevertheless it gave better results than *all* previous converters as well as adequately demonstrating the inherent drift-free performance of digital equipment.

At the time when we produced the experimental model we intentionally chose to fit only three-line interpolation, even though we recognised that it would result in an unfortunate compromise between vertical resolution and the production of 'funnies' on sloping lines. This decision was made for two reasons: (1) it significantly decreased the time taken to obtain operational experience; and (2) suitable Schottky devices were not yet generally available to allow us to realise five-line interpolation in a reasonable space.

It had been our intention to strip this first 'test-bed' converter down and to rebuild it with an

improved performance and to make it bi-directional. However, clearly it would have been illogical to wreck the best 525-to-625 converter then in existence; a more sensible approach was to use it to improve the quality of conversions on Independent Television while building the improved version from scratch. This was the course adopted.

We had to start again, but this did not mean that we had to start from square one. In fact we were restarting from a position of considerable experience and achievement. We could aim very high and we did; we wanted to produce a converter that would be sufficiently close to the theoretically-feasible performance that it would be essentially *ideal;* at the same time we would make it bi-directional.

Some people might consider that we must have had our 'heads in the clouds' to aim so high. If that is the case then indeed we must be very tall – for we had our feet planted firmly on the ground. We knew what we had already achieved and that our goal was not too large a step away.

In this further development we have been successful in substantially meeting our target; but it would be foolhardy to make claims that this equipment converts pictures without impairment; to do so would not only be untrue, but would also insult the reader's intelligence, since how can a 625-line picture with a video bandwidth of 5.5 MHz be converted to 525 lines with 4.2 MHz bandwidth without loss of vertical and horizontal resolution?

However, many observers (both skilled and unskilled) consider that the impairments produced by this converter are 'just perceptible imperceptible' or in other words that the output is rather *better* than the input. This sounds impossible – but there is a good theoretical explanation of this apparent absurdity; certain impairments of the input signal are decreased or removed by the converter; the digital video processing may thus enhance the subjective picture.

To get back to history. We were asked by Independent Television News to make an extra converter whilst we were making the prototype bi-directional machine. Our prototype machine was shown at Montreux in May 1975; at the same time the second two-way converter was delivered to ITN.

In February 1975 an agreement had been signed enabling Marconi Communication Systems Ltd to manufacture and market DICE. A production converter (Fig. 2) made by the company was shown at the National Association of Broadcasters Convention in Chicago in March 1976; others are being produced and have been installed in several countries and we hope this British design will be instrumental in improving communications and understanding between the peoples of the world.

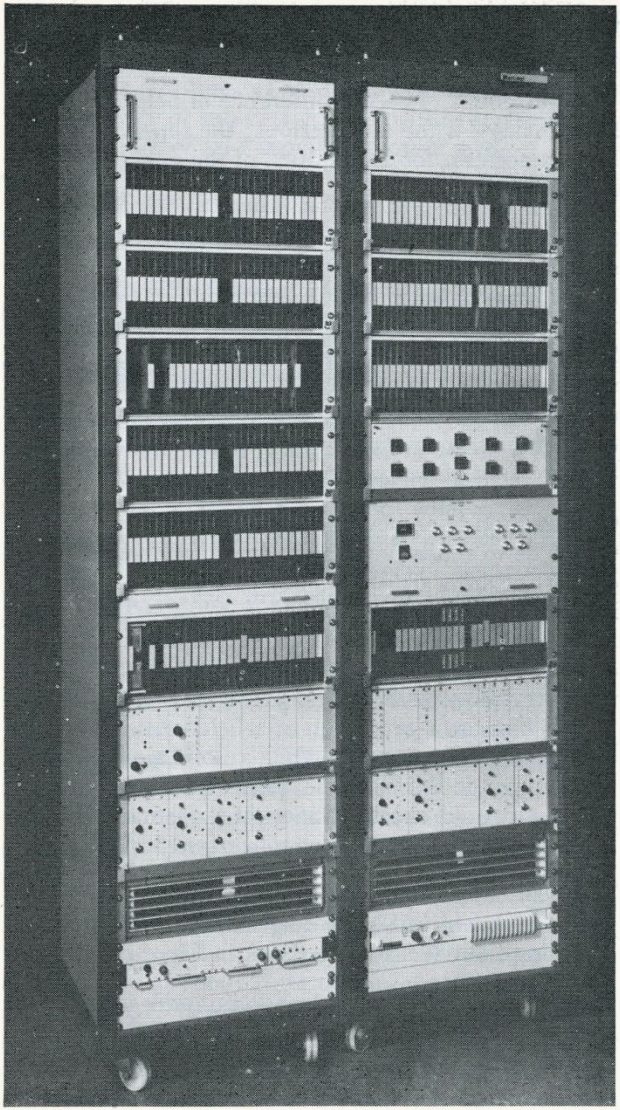

Fig. 2. A production two-way DICE converter, shown with the front panel covers removed. This is typical of the units made to the IBA design by Marconi Communication Systems Ltd and already installed and operational in several countries.

Conversion Principles

The more important differences between the 525 and 625-line standards are:

	525/60	*625/50*
(1) Field frequency	60 Hz	50 Hz
(2) Number of lines	525	625
(3) Field blanking (lines)	21	25
(4) Line Period	63.55μs	64μs
(5) Active Line Period	52.6μs	52μs
(6) Colour coding	NTSC	PAL or SECAM

Without doubt (1) – the difference of field frequencies – is the most serious; the repercussions from it account for about 50% of the entire cost of a reversible converter.

Since there are roughly five 625-line fields for every six 525-line fields, it follows that there should be 1.2 times as much movement between fields of the 625-line picture as occurs between fields of the 525-line picture. Not only *should* there be this relationship but on average there *must* be, since otherwise a one-hour programme would last either fifty minutes or seventy-two minutes, depending on the direction of conversion.

One method of achieving the correct average would be to repeat every fifth input field, when converting from 625-to-525, or to omit every sixth input field when converting from 525-to-625. Although this would result in the programme duration being correct, it would prove unacceptable whenever movement was occurring in the picture. This is because a repetition of a field causes a momentary static picture while an omission causes a momentary doubling of the speed; both of these processes would cause jerkiness of movement of a type normally referred to as movement judder.

The way of effectively neutralising this movement judder is shown in highly diagrammatic form (ignoring the effects of interlace) in Fig. 3.

When the slider (C) of the potentiometer is at the *A* end the output has a timing corresponding to the last field. As the slider progressively moves towards *B* the position of any moving object on the output will progressively become appropriate to an earlier time until, when the slider is at *B*, the position will be appropriate to a time one field earlier.

Suppose we are converting from 525-to-625 and that for one particular output field a timing corresponding to the mid-position of the slider is appropriate ie, the output signal equals $\frac{1}{2}A + \frac{1}{2}B$ or $k = 0.5$. One *input* field later the output information needs to be appropriate for a time one *output* field later. Now one output field corresponds to 60/50 or 1.2 input fields. In the meantime, a new input field has appeared instead of the last input field and the last input field has become the previous input field. Without moving the slider, the output signal has changed to a time appropriate to one input field later; by now moving the slider by one-fifth of the distance from *B* to *A*, (ie, to $k = 0.3$) we add an additional time of a fifth of a field, so that the information at the output has now become appropriate to a time 1.2 input fields later. This process is referred to as *movement interpolation.*

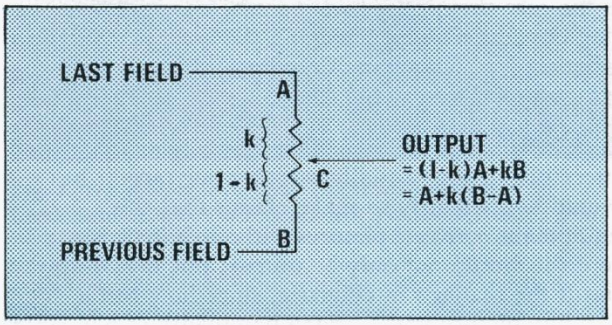

Fig. 3. A highly diagrammatic representation of the way in which the movement judder that might be expected to result from the difference in the number of fields in the 525-line and 625-line television standards is neutralised by providing an output signal containing differing proportions of two successive fields. When the potentiometer slider (C) is at position *A* the output has a timing corresponding to the last field, but as it is moved progressively towards *B* the position of any moving object in the output signal becomes progressively more appropriate to the timing of the previous field.

Whereas in the 525-to-625 direction, the proportion of the later field increases by a fifth (the earlier field decreasing by a fifth for each field); in the 625-to-525 direction the proportions change by a sixth, the later field decreasing by a sixth and the earlier field increasing by a sixth for successive output fields. By this means, the information at the output now changes by steps appropriate to 5/6 of the duration of an input 625 field, ie, by five-sixths of a fiftieth of a second, that is a sixtieth of a second.

In reality, the potentiometer shown in Fig. 3 does not exist but in principle the system would work.

The output from the potentiometer would equal $kB + (1-k)A$, which is equivalent to $A + k(B-A)$.

The video signals at A and B are in eight-digit, pulse-code-modulation (pcm) form. One method of movement interpolation would be to multiply the B signal by k and to multiply the A signal by $(1-k)$ and add the two products together. The variable k can be adequately defined by a six-digit number, as also can $(1-k)$. An eight-by-six digit multiplier would produce a 14-digit answer for the product; in this particular case when the two products are added together the answer could also be completely described by 14 digits. An accuracy of 14 digits (although theoretically achievable) is financially unjustifiable and the aim should be to use only eight bits without causing impairment. If rounding-down to eight digits is adopted (that is, the least significant six digits are ignored), this causes brightness errors which may be equal to the least significant digit of an eight-digit binary word. This is likely to cause a small but visible brightness modulation which will be most noticeable in uniform areas of the picture.

Figure 4 shows how this may be overcome by using what we term 'difference interpolation'. In this system, rounding-down does not result in brightness modulation in uniform areas of the picture but in a very slight departure from the correct amount of interpolation; but only to an imperceptible degree. This is described in more detail in 'Movement interpolation and store control' by Corbyn and Greenfield.

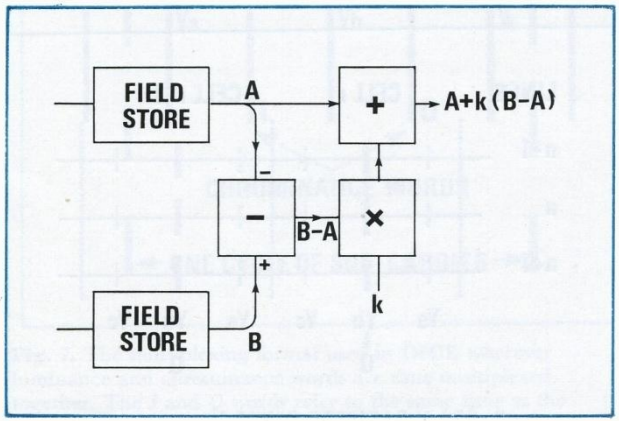

Fig. 4. Rounding-down brightness errors which would arise in practice from the simplified approach outlined in Fig. 3 can be overcome by the technique of 'difference interpolation' shown above. In this system, rounding down of the 14-digit signal that would result from an eight-by-six multiplication of an eight-digit word does not result in any brightness modulation but only in imperceptible departures from ideal interpolation.

Field Stores

Not only has interlace been ignored up to this point but so have many other problems. These will be treated later, but first consider the field stores shown in Fig. 4.

The information when written into the stores is always in a 525-line NTSC form with three words for each cycle of sub-carrier: one word representing the maximum of an I signal; the other two corresponding to angles 120° before and 120° after the peak of the I signal. Each word contains eight digits.

The storage always uses a 525-line standard since storage of 625-line signals would require almost 50% more space. Even for 525-lines, a storage of 2.4 megabits of information is required.

In the 525-to-625-line direction, the received analogue video signal is filtered to remove out-of-band components, and passes to an analogue-to-digital converter (adc) where it is sampled at three times subcarrier frequency. The output words each have eight bits and are written in turn progressively first into one field store and then into the other.

In the 625-to-525 direction the operation is more complex. First of all an analogue PAL or SECAM decoder is used to produce colour signals which are matrixed together to give I and Q signals which individually pass to their own analogue-to-digital converter. A very sharp cut-off filter is used to provide a luminance signal which passes to its own adc.

The outputs from the three adc's are time-multiplexed together and pass through the line interpolation process. Here the number of lines is reduced to 525 per 625-frame and these are interspersed with 100 blank lines which will be discarded. The information from the interpolation is demultiplexed into Y, I and Q words which pass via a digital NTSC coder to the field stores. This coded signal is written into the field stores in the same way as for the 525-to-625 direction, save that the 100 blank lines per frame are ignored.

The information written into the field stores is the same in both directions of operation and has a 525-line NTSC format.

The major need for the field stores is to provide a buffer store of video information, accepting information as it is provided by the source and holding it until it is required for the output signal. The same information is normally used twice for the output signal, once when it represents the last input field and a second time when it represents the previous input field. When information is read from a field store it is also recirculated and written in again, but the normal writing operation destroys previous information.

In a similar (but reverse) way to the 100 blank lines being ignored in the writing into the stores in the 625-to-525 direction, so the read-out from the stores in the 525-to-625 direction is interrupted for 100 lines on each output frame. These 100 lines are nominally spaced equally in time (approximately one for each five normal lines) and are referred to as 'gaps'. The blank lines and the gaps account for the vast majority of the difference in field frequencies between the 525 and 625 standards; most of the remainder is allowed for by the difference between the line frequencies.

Sampling Patterns

On the 525-line side of the machine, irrespective of the direction of conversion, there are two phases of sampling which are used on alternate lines, as shown in Fig. 5. On one line, the plus *I* line, every third sample is along the *I* axis towards plus *I*. On the other line, every third sample is again along the *I* axis but this time towards minus *I*; this is called the 'minus *I* line'. In going from plus *I* to minus *I* the phase has changed by 180°.

It will be recalled that there are 227½ cycles of subcarrier per line, the odd half cycle being used to reduce the visibility of the dot pattern on NTSC in coloured areas. The 180° phase change of sampling neutralises the 'odd' half cycle and causes the samples to line up in vertical columns as though the sampling had been at a multiple of line frequency even though this is *not* the case.

This arrangement of samples is very valuable in obtaining high performance in movement and line interpolation and in the spatial filters described in later sections.

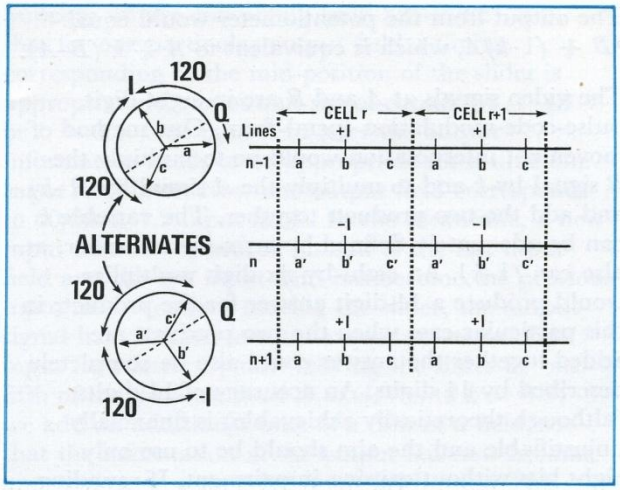

Fig. 5. For the 525-line NTSC system two phases of sampling must be used on alternate lines, irrespective of the direction of conversion. On one line (the 'plus *I* line') every third sample is along the *I* axis towards plus *I*. On the other line, every third sample is again along the *I* axis, but this time towards minus *I*. The 180° phase change neutralises the 'odd' half cycles of subcarrier and causes the samples to line up vertically as though the sampling had been at an exact multiple of line frequency, even though this is *not* the case.

On the 625-line side of the machine, irrespective of the direction of conversion, the sampling occurs at an integral multiple of line frequency. Again this causes the samples to line up in vertical columns as shown in Fig. 6. The sampling frequency on the 625

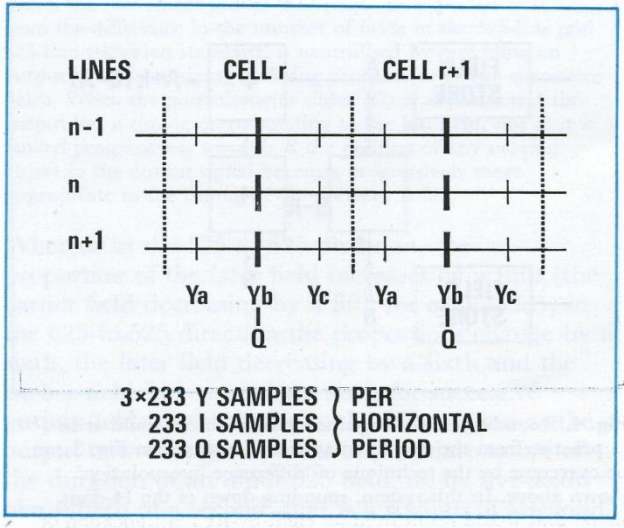

Fig. 6. On 625-lines, irrespective of the direction of conversion, the sampling occurs at an integral multiple of line frequency, again causing the samples to line up in vertical columns.

side of the converter changes slightly, causing the number of groups of three samples (a cell) to change from 231 in the 525-to-625 direction to 233 in the 625-to-525 direction. The main reason to change from 227½ to (nominally) 232 is to compensate for the change in the ratio of the active picture time to the line period. The use of 231 in one direction of operation, and 233 in the other, causes a very small horizontal stretch of the picture so that, with nominal signals, a trace of the picture is blanked out, so defining the blanking time.

Multiplexed Format

In the converter wherever luminance and chrominance words are time multiplexed together a standard has been adopted, as shown in Fig. 7. In a cell, which lasts a nominal cycle of subcarrier, three luminance words occur: Y_a, Y_b, Y_c. In between the first two an I word is multiplexed; in between the second two a Q word is inserted. It is arranged that the I and Q words refer to the same time as the middle luminance word, Y_b.

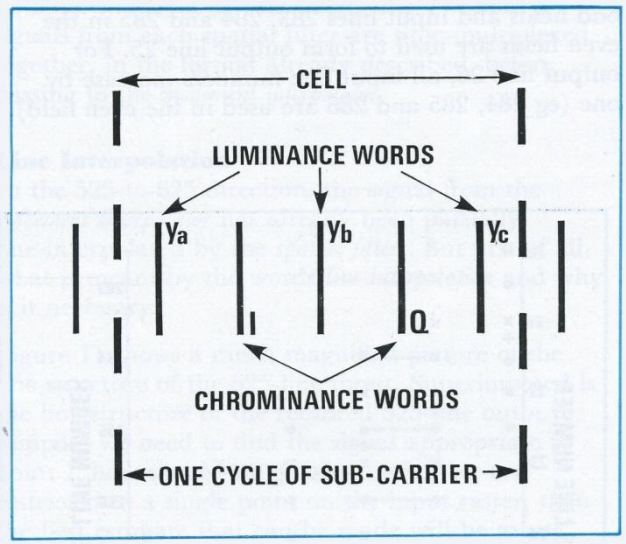

Fig. 7. The multiplexing format used in DICE wherever luminance and chrominance words are time multiplexed together. The I and Q words refer to the same time as the middle luminance word, Y_b.

Effects of Colour Coding and Interlace

In describing movement interpolation it has already been shown that two successive fields have to be added together in different proportions; to understand line interpolation it needs to be shown how five successive lines are, in an algebraic sense, added together.

For an NTSC signal, the phase of subcarrier shifts on consecutive lines by 180°; imagine what would happen if we were to take the arithmetic mean of the signal on two consecutive lines in order to obtain a signal approximating to what would have been obtained if we had scanned the original scene mid-way between these two lines. When these two lines were averaged, the luminance would average correctly but, due to the change of phase of 180° for the subcarrier, the chrominance would cancel out in uniformly colour areas. It is thus essential to alter the coding system to remove this 180° phase change, or better still to decode the chrominance into base-band signals *before* averaging. The separation of luminance and chrominance and the decoding of chrominance into I and Q components are performed in a spatial filter, so called because it operates in two dimensional space, unlike ordinary filters which, effectively, are one dimensional. The design of these special filters is described in 'Spatial filters' by Baldwin and Thirlwall.

There is one major converter subsystem, shown in Fig. 8, which remains substantially unchanged by the direction of operation of the converter; this

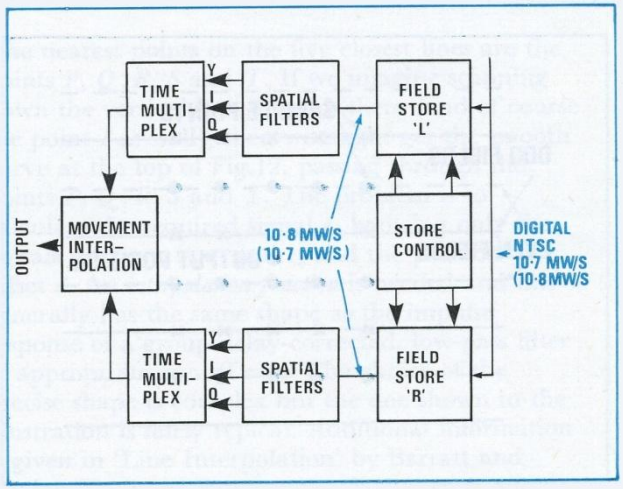

Fig. 8. The whole of this major subsystem, including the two field stores, remains substantially unchanged for both 525-625 and 625-525 directions of conversion. The two field stores are each connected via spatial filters and time multiplexers to one input of the movement interpolator.

consists of the two field stores each of which is connected via a spatial filter and a multiplexer to one input of the movement interpolator.

In the 525-to-625 direction the converter input signal (after being sampled at three times NTSC subcarrier frequency, approximately 10.7 MHz) is converted to eight-bit words and is written into one of the field stores during one field, and then into the other, the two field stores being 'written' during alternate fields. Simultaneous 'reading' of both field stores occurs at an appropriate time to provide the output signal, making allowance for the delays in the spatial filters and elsewhere. However, reading of the stores occurs at a greater speed, approximately 10.8 Mw/s, to shrink the active line time in the correct proportions.

It helps to clarify this process by not thinking of the information in the stores as a television signal. It is essentially different in that the information describing each picture is permanently there, in the store, until replaced by more recent information. It is rather like an analogue of the original *scene* and has to be scanned to produce a normal television signal. The problems of colour coding and of interlace are overcome by the spatial filters whose principle of operation is shown in Fig. 9. Each spatial filter considers information appearing at five points on each of three successive lines of a field. One operates on odd fields; the other operates on even fields.

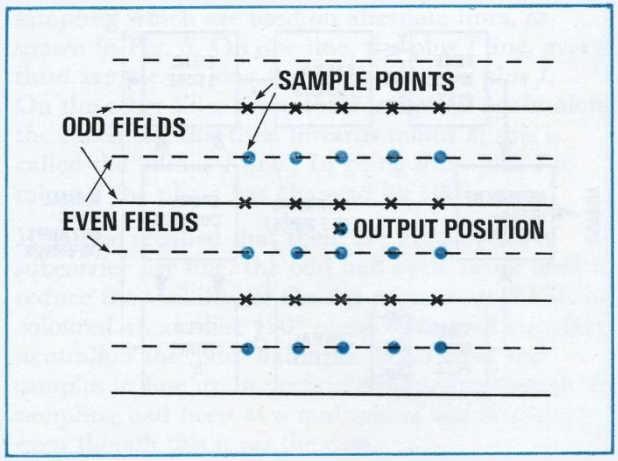

Fig. 9. The spatial filters each take different proportions of 15 samples, representing five points on each of three successive lines of a field. One spatial filter operates on odd fields; the other on even fields. In this diagram, for odd fields, the sampling points are represented by small crosses.

For the odd fields, the points are shown by crosses. The signals appearing at these points are combined in certain proportions to yield a luminance signal appropriate to the output position indicated, and in different proportions to give a chrominance signal appropriate to the same output position. In addition, the chrominance is separated to yield I and Q components which are both demodulated.

The other field is simultaneously treated by a similar filter to give Y, I and Q components for the same output position, hence neutralising the effects of interlace. As successive words arrive, the information progressively steps along the five points in each line; so that the output word progressively represents the Y, I and Q components appearing along one line (a line, incidentally, which is half way between the lines of odd and even fields).

Although this happens for some output lines, the operation passes through a sequence, part of which is shown in Fig. 10. This shows which input 525 lines are used to form each output line. The description so far indicates how input lines 20, 21 and 22 in the odd fields and input lines 283, 284 and 285 in the even fields are used to form output line 25. For output line 26, all input line numbers increase by one (eg 284, 285 and 286 are used in the even field).

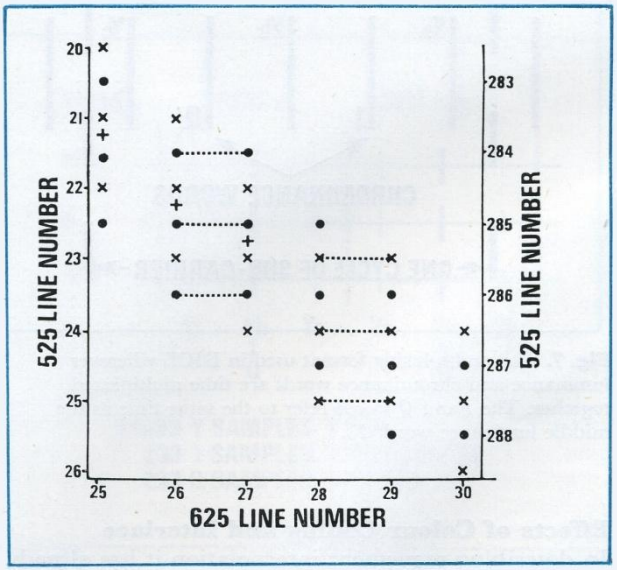

Fig. 10. Showing which input 525-lines are used to form each output line, in a repeating sequence.

However, for output line 27 the same input lines in the even field are used again, this occurring when a 'gap' (no information) comes from the store holding the even field. Three, one-line-long, shift-registers in the spatial filters recirculate to make this information available a second time.

Prior to this line, the 'even' spatial filters will have been giving an output position a quarter of the pitch of the lines of a field above the middle input line; now it changes to below the middle input line. The 'odd' spatial filter changes from a quarter of a line-pitch down, to a quarter-pitch up.

By this means, half the line-interpolation has been provided in the spatial filters and the 'gaps' from the field stores have been filled.

In the 625-to-525 direction, the 'gaps' do not occur and the spatial filters shift the output in the same direction throughout a whole field but reverse direction in alternate fields; this has the effect of restoring interlace.

In both directions of conversion, the Y, I and Q signals from each spatial filter are time-multiplexed together, in the format already described, before passing to the *movement interpolator*.

Line Interpolation

In the 525-to-625 direction, the signal from the *movement interpolator* has already been partially line-interpolated by the *spatial filters*. But first of all, what is meant by the words *line interpolation* and why is it necessary?

Figure 11 shows a much magnified picture of the line structure of the 625-line input. Superimposed is the line structure of the required 525-line output. Suppose we need to find the signal appropriate to point I, how should we go about it? If we are restricted to a single point on the input raster, then the best estimate that can be made will be to use the signal-level appearing at the nearest point, that is point R. Inevitably, this point must appear along the vertical line through I; in this instance R appears below I. Since I is about 40% of the distance from R to Q, a better estimate would be obtained by adding 40% of the signal at Q to 60% of the signal at R. If two points are better than one, then it follows that three points would give an even better estimate, and so on. Whilst this is theoretically true, the improvement in accuracy of the estimate, as the number of points increases, very rapidly slows down. The first point is 'essential', the second 'very important', the third 'most advantageous', the fourth 'desirable'. By the time the fifth point is reached it could be described only as 'beneficial'. To increase beyond this point would be virtually useless, since the improvement would at best be detectable only on specially selected test patterns.

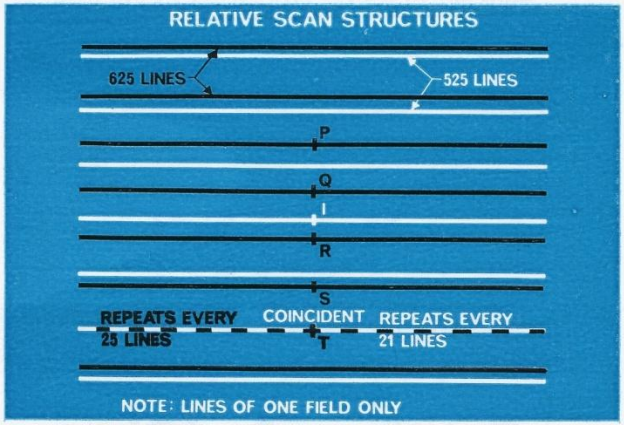

Fig. 11. Much magnified illustration of the line structure of the 625-line input on which has been superimposed the line structure of the required 525-line output. While theoretically it might appear that the more points that are used to determine the output signal, the greater the accuracy, it can be shown that in practice to extend this to beyond five points would be virtually useless since any improvement would be detectable only on specially selected test patterns.

The nearest points on the five closest lines are the points P, Q, R, S and T. If we imagine scanning down the vertical line through them (and of course the point I as well), then we might get the smooth curve at the top of Fig.12, passing through the points P, Q, R, S and T. The problem is to calculate the required signal e_d knowing only the signals e_{-2}, e_{-1}, e_0, e_1 and e_2 and the position offset d. An *interpolation function* is needed and this generally has the same shape as the impulse response of a group-delay-corrected, low-pass filter of appropriate cut-off rate. The choice of the precise shape is complex but the one shown in the illustration is fairly typical. Additional information is given in 'Line Interpolation' by Barratt and Taylor. To interpolate, centre the interpolation function at the required output position, read off the ordinates a_{-2}, a_{-1}, a_0, a_1 and a_2 in line with each input point, multiply each of these by the signal appearing at that point and add together the

results of those multiplications. This yields the required signal. One characteristic of the interpolation function is that the sum of the intercepts a_{-2}, a_{-1}, etc, must be constant and independent of the offset d. If this were not the case, a uniform grey picture after interpolation would yield a grey signal modulated by the variation of the sum, as the offset d varied from line-to-line.

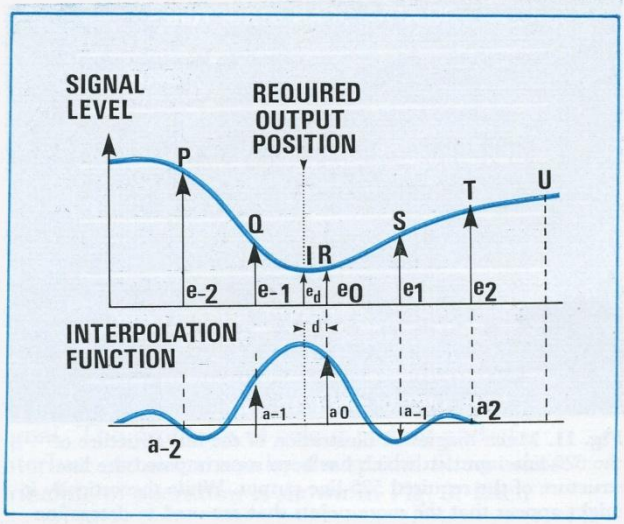

Fig. 12. Five-line interpolation. The top curve represents signal levels at points *P, Q, R, S, T* on five lines in the form of a smooth curve. To calculate the required signal from the known signals and the offset '*d*' an 'interpolation function' is needed. While the choice of the precise shape of such a function is complex that shown in the lower part of the illustration is fairly typical.

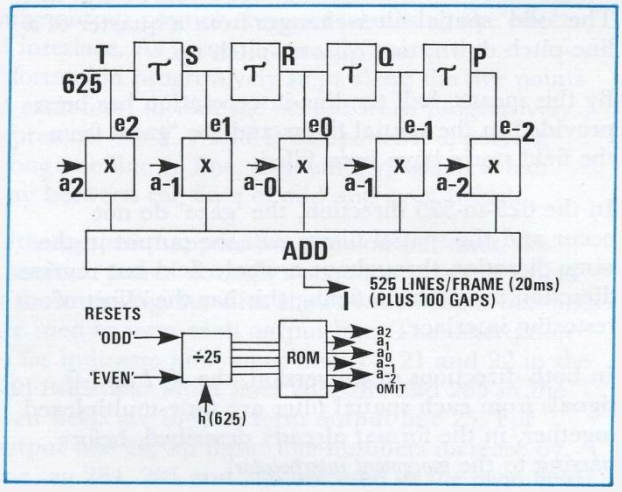

Fig. 13. One approach to the implementation of a five-line interpolator. This system, however, would tend to suffer from 'rounding-off errors' which may cause visible brightness modulation in uniform areas of the picture. It should be noted that while the illustration shows only 625-to-525 conversion, the only significant difference for the 526-to-625 direction would be in the generation and size of the coefficients.

Figure 13 shows a block diagram of how such a 5-line interpolator can be realised; the description is for 625-to-525 conversion but the only significant difference is in the generation of the coefficients and particularly the size of the coefficients.

First of all we must generate the coefficients equivalent to the intercepts already mentioned. Since 625 and 525 are both divisible by 25 there must be 25 repetitions of the interpolation per television frame, each repetition lasting 25 and 21 lines respectively for the two line standards. The 625 horizontal frequency is used to clock a counter which divides by 25. This counter is reset to one particular number at the start of an 'odd' field; to a different number at the start of an even field. The state of the counter provides the address used to interrogate the read-only-memory which stores the coefficients required for interpolation. Four times during the sequence, the magnitude of the offset d becomes greater than it would one input-line later, and when this occurs an 'omit' pulse is generated. This ensures that the best group of five lines is always used. Without it, we would normally be extrapolating rather than interpolating. This problem does not arise in the 525-to-625 direction, nor do 'gaps' occur. The four labelled blocks at the top of the illustration use shift registers to provide a delay of precisely one input line per box, so that the outputs of these blocks, together with the input of the first line, provide the same point on five successive input lines. Because the first line has to be delayed more than the second, and so on, the points *T, S, R, Q* and *P* appear in the reverse order. The signal appearing at each point is multiplied by its appropriate coefficients and the products are added together to yield 525 lines per input frame, interspersed with one hundred 'gaps'. From previous experience, however, we had learned that this particular arrangement has a number of undesirable characteristics; for example rounding-off errors are difficult to control, tending to cause visible noise in uniform areas of the picture.

Figure 14 shows how these problems are overcome.

Instead of using the signals at P, Q, S and T directly, difference signals $\triangle_{-2}$, $\triangle_{-1}$, $\triangle_1$ and $\triangle_2$ are produced by subtracting the signal e_0 at R. These four delta signals are then multiplied by four of the coefficients obtained from the interpolation function, as described before, the centre one being unused. The sum of the products is added to the signal e_0 appearing at R, and providing the sum of the coefficients including the central one totals one, the interpolation produced is identical to the previous system at least in theory.

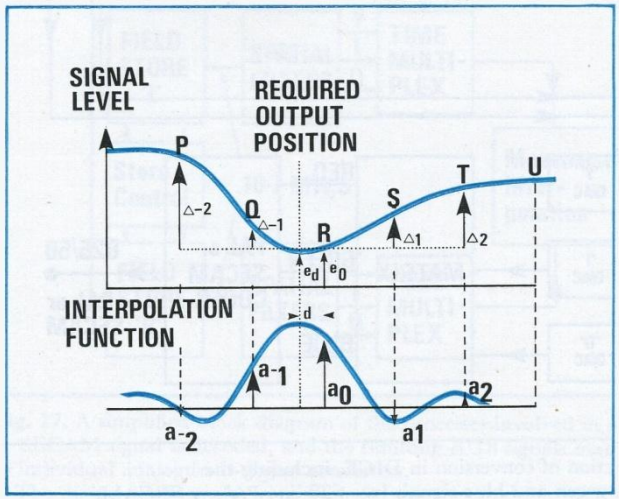

Fig. 14. A preferred approach to five-line interpolation using 'difference' rather than 'direct' signals. As indicated earlier (Fig. 4) this approach successfully overcomes the problem of rounding-off errors.

Figure 15 shows the block diagram of the 5-line 'difference interpolation' used in the converter, the essential differences being the addition of four subtractors to provide the delta signals and the removal of the central multiplier. The coefficients can be generated as before, but now, of course, only four are required. Some of the important advantages gained by adopting this difference interpolation system are:

(1) In the more uniform parts of the picture, the rounding off errors become vanishingly small.

(2) Errors in this part of the converter tend to affect the accuracy of interpolation when difference interpolation is used, rather than appearing directly as changes of brightness. For example, if one of the outer multipliers be removed, the accuracy drops to about the equivalent of four-line interpolation, whereas, for the system described earlier, it would cause intolerable line-to-line brightness variations.

As noted earlier, the sampling pattern on every line is the same and the time multiplexing is constant. This means that although the one interpolation system operates on all the Y, I and Q signals it does so one at a time. For example, if one of the points T, S, R, Q or P is an I signal then all of them are I signals. This results in the same multiplexed format coming from the interpolator as was applied to the input.

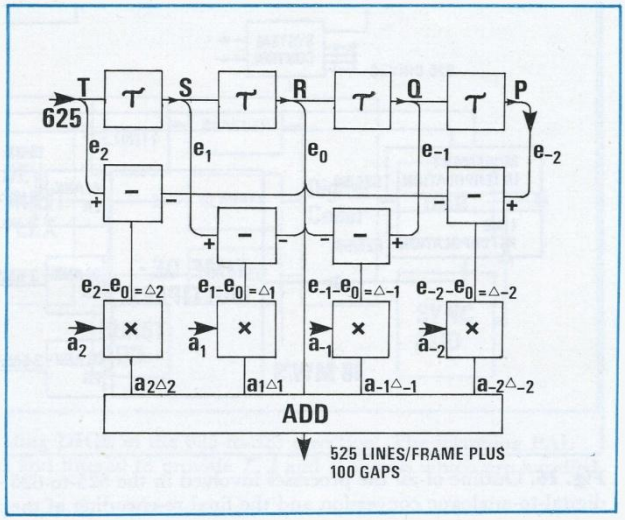

Fig. 15. Block diagram of the five-line difference interpolation used in DICE. This differs essentially from the system outlined in Fig. 13 by the addition of four digital subtractors to provide the delta signal and the removal of the central multiplier. The coefficients are generated as before but now only four are required. Such an arrangement not only minimises any rounding-off errors but also reduces the significance of any other errors in this part of the converter.

We are now almost at the end of the processes shown in the 525-to-625 block diagram (Fig.16). Conversion is now complete; all that remains to be done is to demultiplex the signal to yield Y, I and Q words which pass through their appropriate blanker to their *digital-to-analogue converter*. The outputs of the dac's are filtered, corrected for $(\sin x)/x$ loss, and matrixed together to give red, green and blue signals which pass to a conventional analogue PAL or SECAM coder.

In the 625-to-525 block diagram (Fig.17) the majority of processes have already been described.

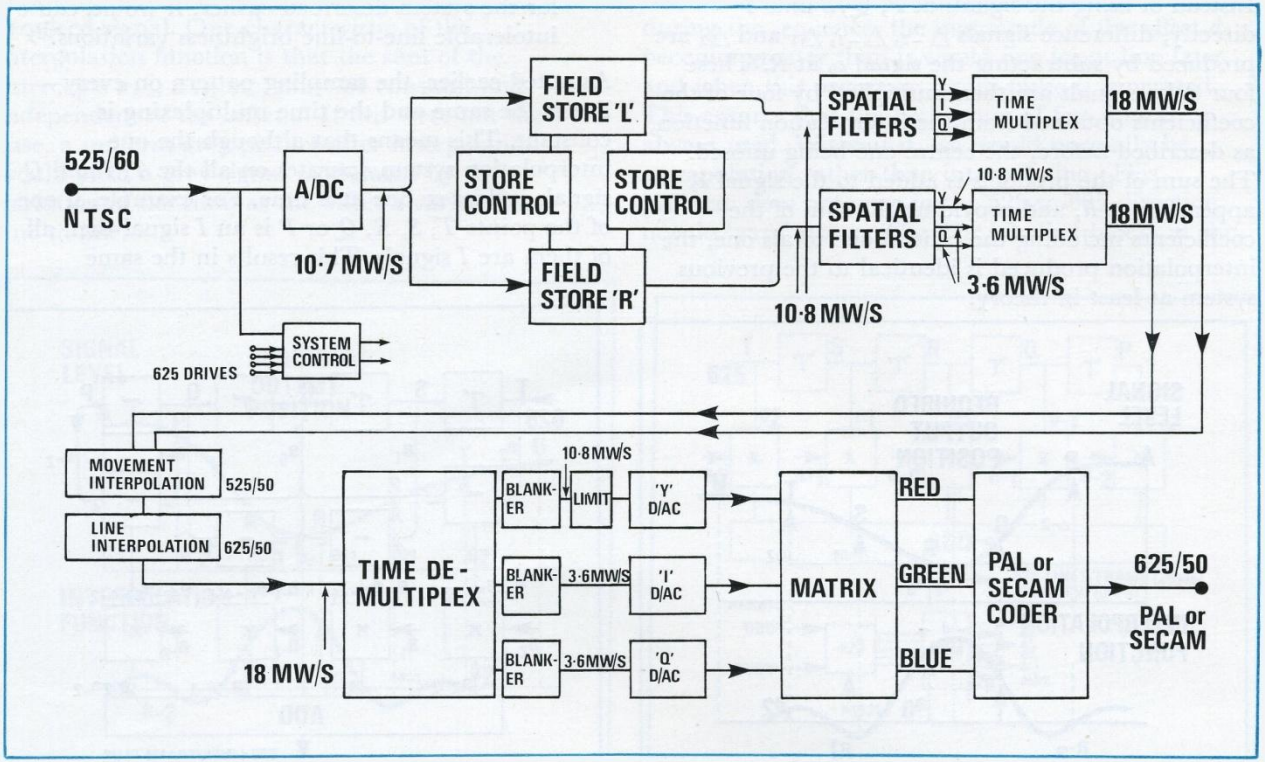

Fig. 16. Outline of all the processes involved in the 525-to-625 direction of conversion in DICE, including the digital-to-analogue conversion and the final re-encoding of the red, green and blue signals into 625-line PAL or SECAM.

The incoming PAL or SECAM signal is decoded and the resulting red, green and blue signals matrixed to give I and Q signals. The incoming signal is also carefully filtered to yield the luminance or Y signal. The Y, I and Q signals pass to their individual analogue-to-digital converters. (Further information is given in 'Analogue processing and facilities' by Lever and Connolly.) The words produced by the adc's are time multiplexed together and pass to the *line interpolation*. The signal is then demultiplexed and turned directly into an NTSC type of signal in digital form; as explained in 'Digital coding and blanking' by Carmen and Bellis. The coded signal then passes through the field stores, spatial filters and time multiplexers to the movement interpolator as already described. The movement interpolated signal is demultiplexed, the luminance signal limited and burst forming pulses added to the I and Q signals and each signal carefully blanked; the I and Q blankers operate once to shape the leading and trailing edges of the burst and a second time to conventionally blank the chrominance picture information. The Y, I and Q signals pass to the digital coder which produces a standard NTSC signal, albeit in digital form and without syncs. This is then converted to analogue form, filtered, $(\sin x)/x$ corrected and syncs added. The process is complete.

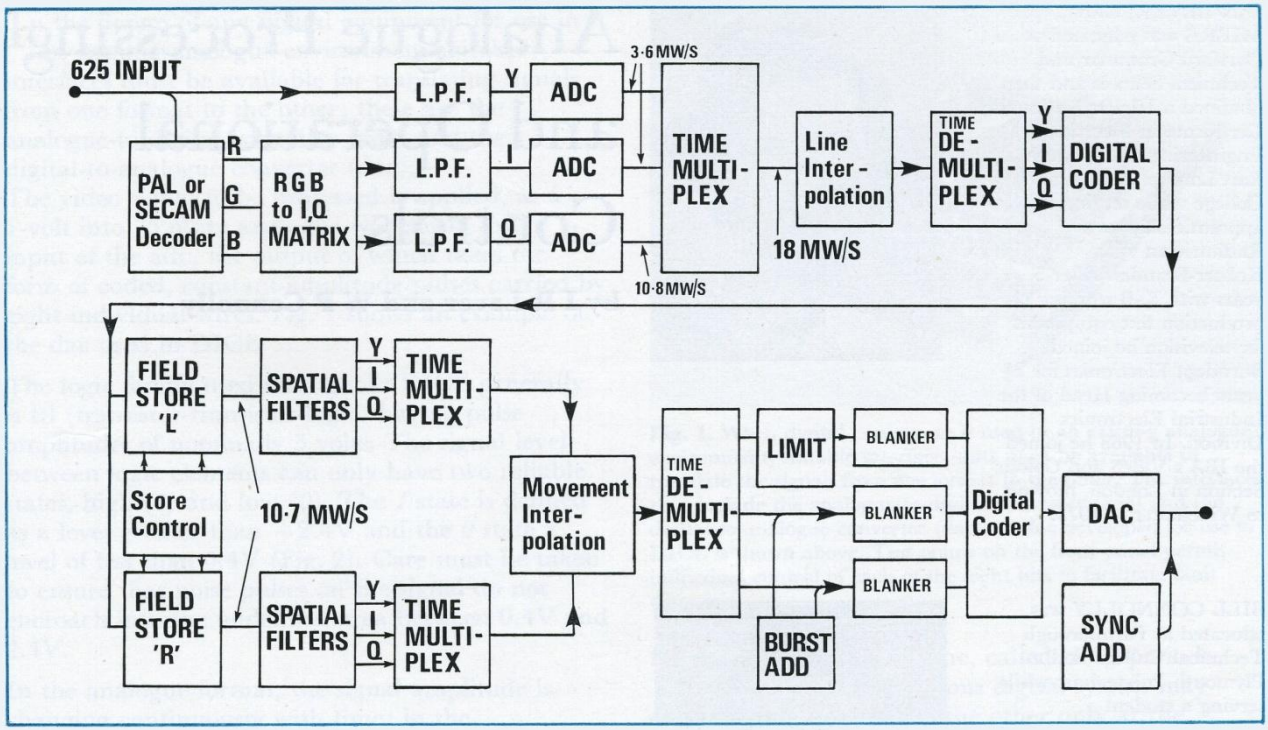

Fig. 17. A simplified block diagram of the processes involved in operating DICE in the 625-to-525 direction. The incoming PAL or SECAM signal is decoded, and the resulting RGB signals matrixed and filtered to provide Y, I and Q signals which are handled in individual analogue-to-digital converters.

The digital words are time multiplexed together and pass to line interpolation after which the signal is demultiplexed and turned directly into an NTSC type of signal but in digital form. The coded signal then passes through the field stores, spatial filters and time multiplexers to the movement interpolator. The signal is then demultiplexed and further signal processing is carried out. A digital coder then produces a standard NTSC signal in digital form and without syncs. This is converted to analogue form, processed and syncs added to provide a conventional 525-line NTSC output signal.

Analogue Processing and Operational Controls

by I R Lever and W P Connolly

IAN LEVER, CEng, MIERE was educated at Dartford Grammar and Technical Schools and then obtained a Higher National Certificate in Electrical Engineering at the South East London Technical College while serving an apprenticeship as a Radiotrician with Kolster-Brandes. After 5 years with K-B working on production test equipment for television he joined Burndept Electronics for 2½ years becoming Head of the Industrial Electronics Division. In 1968 he joined the IBA's Video and Colour Section in London, moving to Winchester in 1973.

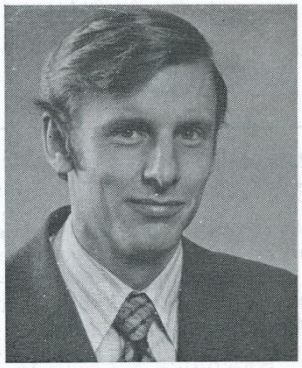

BILL CONNOLLY was educated at Farnborough Technical College and at Plymouth Polytechnic while serving a student apprenticeship with the Ministry of Defence, later working on a variety of military communication and radio relay systems. In 1973 he joined the IBA's Video and Colour Section and since then has worked on DICE and other aspects of digital video signal processing.

Synopsis

The first part of this article describes the interfaces used between the established analogue TV studio world and the growing digital one, with particular reference to their use within DICE. The second part of the article describes the various operational facilities provided by the equipment and their control.

DICE is a digital television standard converter designed to convert pictures in either direction between 525-line 60-field NTSC and 625-line 50-field PAL or SECAM. The most significant advantage to be gained by using digital techniques in the handling of television is the considerable amount of processing that may be applied to the signal without noticeably degrading the picture by the type of impairment inherent in analogue techniques; differential phase/gain, non-linearity, falling signal/noise ratio, etc. As digital circuits cannot drift and do not require adjustment the consistency of performance does not deteriorate with time.

To allow DICE to be integrated into the conventional studio systems, the input and output signal format is the traditional analogue 1-volt into 75 ohms, but within DICE the signal is conveyed as a series of constant amplitude pulses along 8 parallel wires at rates up to 18 million pulses per second on each wire.

In the design of any digital equipment for use in an existing analogue environment suitable interfaces must be available for translating signals from one format to the other; these are the analogue-to-digital converter (adc) and the digital-to-analogue converter (dac).
The video signal to be processed is applied, as a 1-volt into 75 ohms analogue waveform, to the input of the adc, the output of which takes the form of coded, constant-amplitude pulses carried by eight individual wires. Fig. 1 shows an example of the dac used in DICE.

The logic family used throughout DICE generally is ttl (transistor-transistor-logic) having pulse amplitudes of nominally 5 volts. The signal level between logic elements can only have two reliable states, high (*1*) and low (*0*). The *1* state is defined as a level greater than +2.4V and the *0* state a level of less than 0.4V (Fig. 2). Care must be taken to ensure that noise pulses on the signal do not encroach into the undefined area between 0.4V and 2.4V.

In the analogue format, the signal amplitude is changing continuously with time; in the synchronous digital format the signal is either *0* or *1* for specific intervals of time, called clock periods. The signals in a synchronous digital system may change from one level to the other only at the boundaries between these clock periods which, for

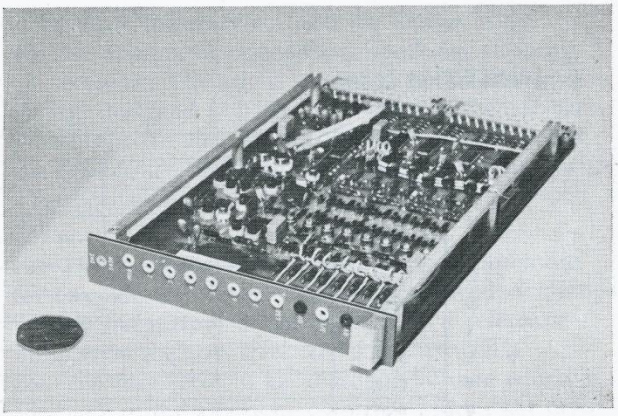

Fig. 1. When digital equipment is used in an existing analogue environment, suitable interface units must be available to translate the signals from one format to the other. The interfaces thus include the analogue-to-digital converter (adc) and the digital-to-analogue converter (dac). A dac developed for use in DICE is shown above. Test points on the front panel permit individual control of each of the eight bits to facilitate fault diagnosis.

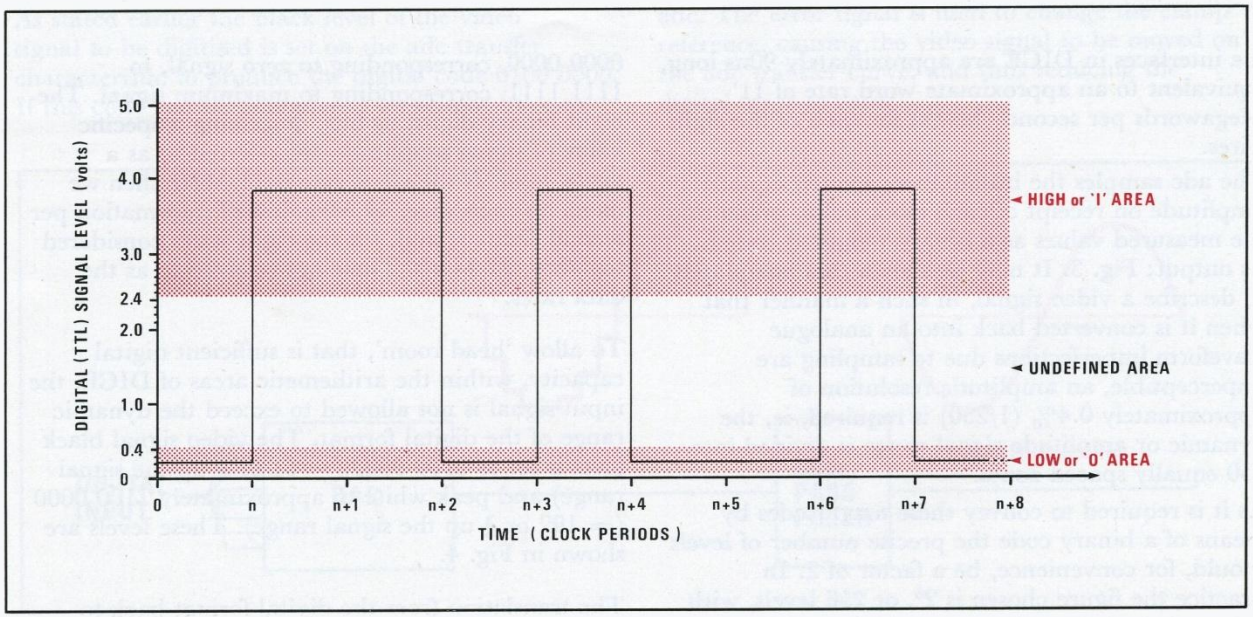

Fig. 2. Analogue signal levels may vary continuously with time but in synchronous digital systems signals can change only at the boundaries of the clock periods. With transistor-transistor-logic (ttl) there are only two states: high '1' and low '0'. The 1 state is defined as a voltage greater than 2.4 V and the 0 state as a level of less than 0.4 V. Care must be taken to ensure that noise pulses on the signal do not encroach into the undefined area between these two levels.

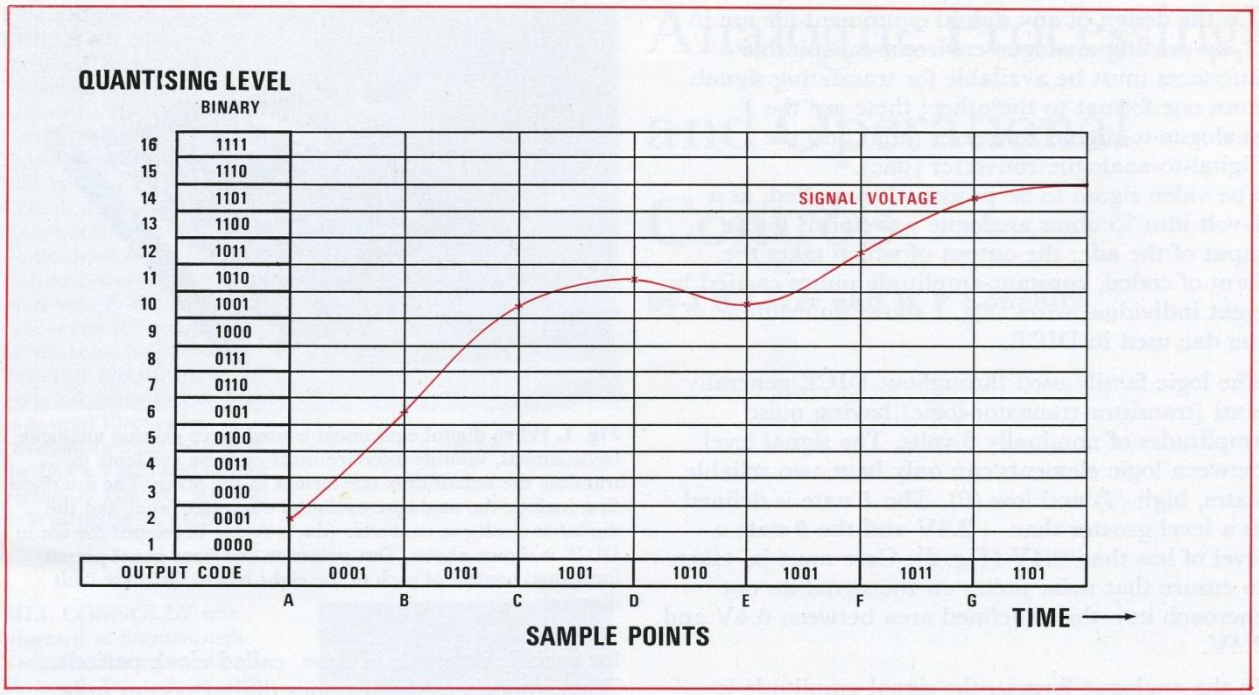

Fig. 3. The analogue-to-digital converter samples the instantaneous video signal amplitude on receipt of each clock pulse, presenting the measured values as a binary-weighted-code at its output. A simple example of a 16-level (4-bit) adc is shown in this illustration.

the interfaces in DICE are approximately 90ns long, equivalent to an approximate word rate of 11 Megawords per second (Mw/s) for each of the eight wires.

The adc samples the instantaneous video signal amplitude on receipt of each clock pulse, presenting the measured values as a binary weighted code at its output: Fig. 3. It may be shown that adequately to describe a video signal, in such a manner that when it is converted back into an analogue waveform imperfections due to sampling are imperceptible, an amplitude resolution of approximately 0.4% (1/250) is required, ie, the dynamic or amplitude signal range is divided into 250 equally spaced zones.

As it is required to convey these magnitudes by means of a binary code the precise number of levels should, for convenience, be a factor of 2. In practice the figure chosen is 2^8, or 256 levels, with the resultant information carried in parallel form on eight wires.

Hence any signal passing through DICE may have only 256 discrete values ranging from 0000 0000, corresponding to zero signal, to 1111 1111, corresponding to maximum signal. The eight binary digits, or bits, describing a specific value of signal amplitude are referred to as a *digital word*. If the clock rate is 11 MHz then we have, on each wire, 11 Megabits of information per second (Mb/s); and, for the eight wires considered together, 11 Mw/s. These are referred to as the data rates.

To allow 'head room', that is sufficient digital capacity, within the arithemetic areas of DICE, the input signal is not allowed to exceed the dynamic range of the digital format. The video signal black level is set to 0100 0000 (= 64 or $\frac{1}{4}$ up the signal range) and peak white to approximately 1100 0000 (= 192 or $\frac{3}{4}$ up the signal range). These levels are shown in Fig. 4.

The translation from the digital format back to the analogue video signal is performed by the digital-to-analogue converter (dac). This unit accepts the digital data and reconstructs the analogue waveform using each data word to define

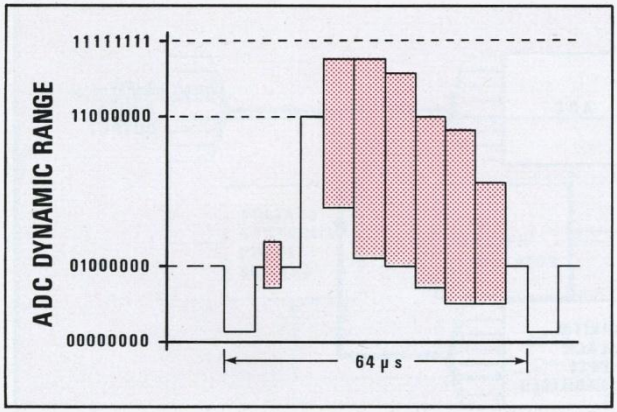

Fig. 4. To allow sufficient digital capacity within the arithmetic areas of DICE, in the presence of a 100% colour bar input, the analogue signal presented to an adc is restricted to less than the complete dynamic range of the digital format.

the instantaneous signal amplitude at each clock pulse. The amplitude of this analogue waveform remains constant for the clock period, hence the output from the dac has a stepped form (Fig. 5). When this signal is passed through a low pass filter a close approximation to the original signal waveform is recovered.

As stated earlier the black level of the video signal to be digitised is set on the adc transfer characteristic to produce the digital code 0100 0000. If this code can be guaranteed we do not have to store black information, and the design of the various arithmetic operations, blanking, clipping, etc, is eased. The adc is preceded by an analogue clamp, designed to hold black level at the correct potential with very low drift. However, the adc is itself liable to slight drift with time and temperature. Therefore steps must be taken to correct any such drift.

For DICE, the technique adopted to compensate for any adc errors was to introduce a feedback loop. The digital output is compared with a reference number in a *digital black level stabiliser* which converts any error to an analogue voltage which is fed back to the analogue clamp to reduce the error to zero: Fig. 6.

The digital output from the adc is monitored during the back porch when the signal is at black. This is the area occupied by the colour burst and suitable processing is required to deduce the value of black from the samples describing the burst waveform. This value is compared with the reference value 0100 0000 and any error is accumulated digitally, the accumulation being the digital equivalent of integration. The output of the accumulator drives a simple dac yielding an error voltage that is fed back to the clamp preceding the adc. The error signal is used to change the clamp reference, causing the video signal to be moved on the adc transfer curve, and thus reducing the digital error to zero.

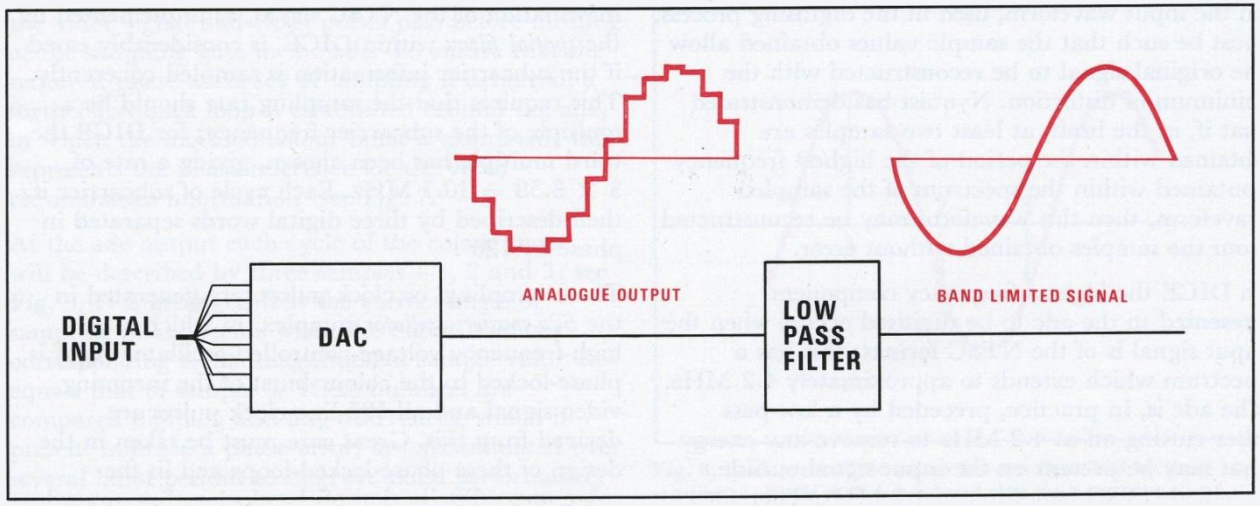

Fig. 5. The digital-to-analogue converter (dac) reconstructs the analogue waveform in a series of steps, using each data word to define the instantaneous signal amplitude at each clock pulse. Subsequently passing the signal through a low-pass filter recovers the original signal waveform with calculable accuracy.

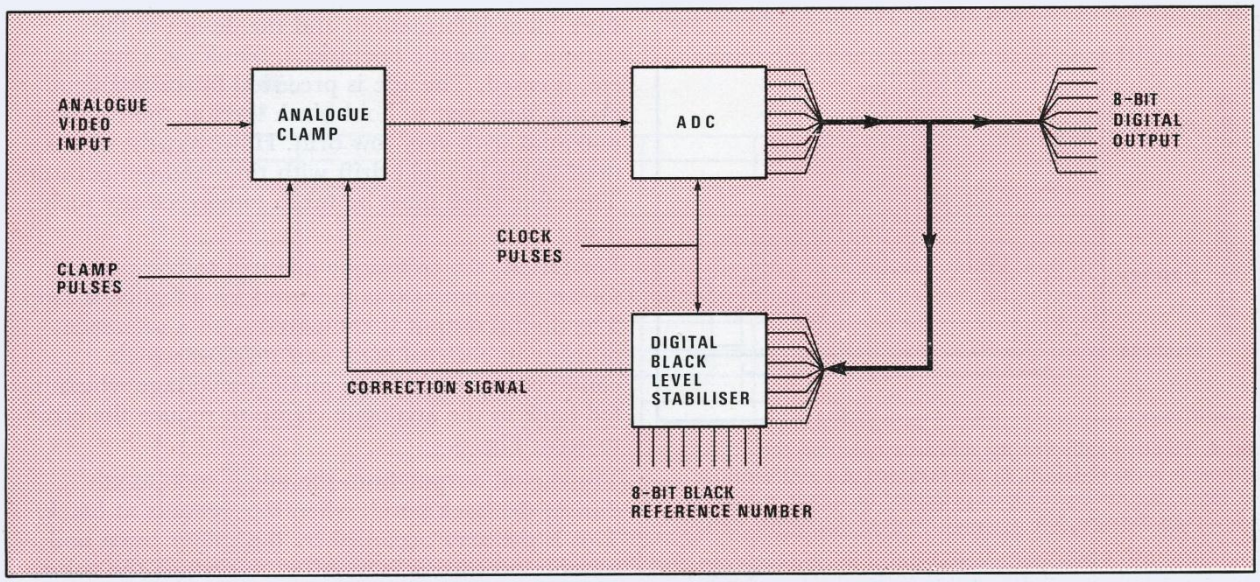

Fig. 6. To simplify arithmetic processing, the picture black level is defined through DICE and is accurately maintained by a digital black level stabiliser forming a feedback loop around the adc. The digital number corresponding to the back porch of the television waveform is computed and compared with an eight-bit reference number: any error is converted to an analogue correction signal which moves the incoming signal up or down the dynamic range of the adc, as appropriate.

If the reference value fed to the digital black level stabiliser is changed, the system will readjust to set video black level to the new number. This facility is used in DICE when 'lift' is required to be removed or added to a video signal while being converted.

The time intervals between the sampling points on the input waveform, used in the digitising process, must be such that the sample values obtained allow the original signal to be reconstructed with the minimum of distortion. Nyquist has demonstrated that if, in the limit, at least two samples are obtained within the period of the highest frequency contained within the spectrum of the sampled waveform, then this waveform may be reconstructed from the samples obtained without error.

In DICE the highest frequency component presented to the adc to be digitised occurs when the input signal is of the NTSC format; this has a spectrum which extends to approximately 4.2 MHz. The adc is, in practice, preceded by a low-pass filter cutting off at 4.2 MHz to remove any energy that may be present on the input signal outside specified NTSC bandlimit of 4.2 MHz. Therefore, according to Nyquist, a sampling rate of 8.4 MHz would be high enough to adequately describe this signal. However, if 8.4 MHz were used the subcarrier components of the NTSC signal occurring at 3.58 MHz would be sampled asynchronously, each sample occurring at a different phase of the subcarrier with respect to burst.

The separation of the luminance and chrominance information of the NTSC signal, as implemented by the *spatial filters* within DICE, is considerably eased if the subcarrier information is sampled coherently. This requires that the sampling rate should be a multiple of the subcarrier frequency; for DICE the third multiple has been chosen, giving a rate of $3 \times 3.58 = 10.7$ MHz. Each cycle of subcarrier is then described by three digital words separated in phase by 120°.

These sampling, or clock pulses, are generated in the *525 master oscillator* complex, in which a high-frequency voltage-controlled-oscillator (vco) is phase-locked to the colour-burst of the incoming video signal and all 525-line clock pulses are derived from this. Great care must be taken in the design of these phase-locked-loops and in the subsequent distribution of the clock pulses around the converter. If any jitter is present or is induced into the clocks the samples derived by the adc will

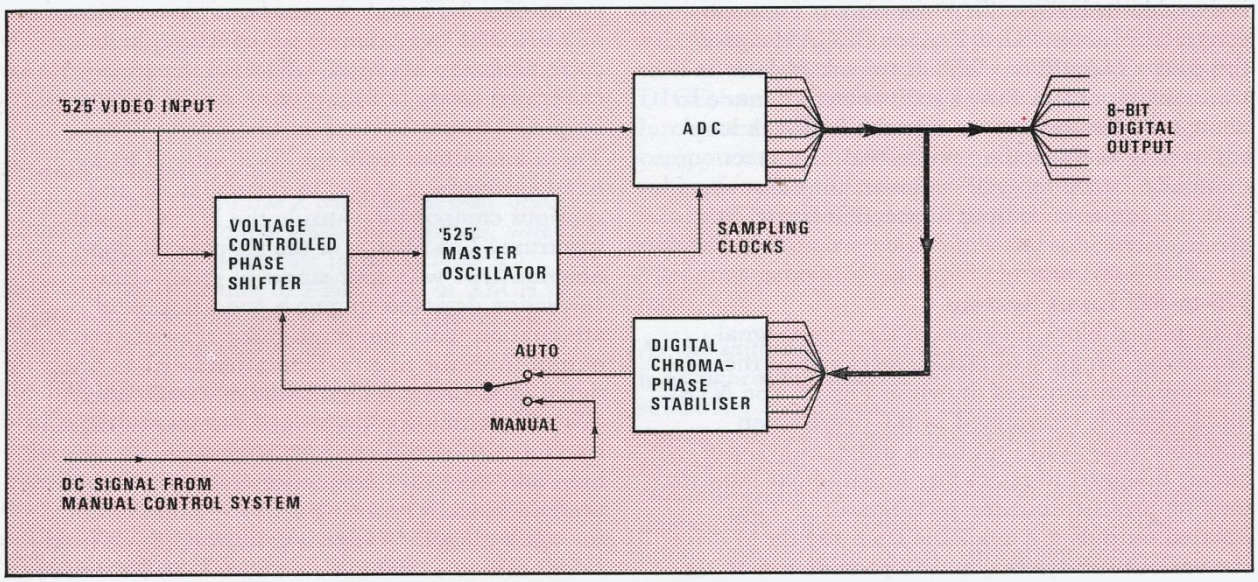

Fig. 7. In order to define the absolute phase of sampling at the adc and so permit digital separation of the 525-line luminance and chrominance information, a digital chroma phase stabiliser examines the adc output during the colour burst period and adjusts a voltage-controlled phase-shifter in the reference video feed to the 525 master oscillator.

not be exactly 120° apart in phase, reducing the efficiency of the spatial filters.

The spatial filters also require that the absolute phase of sampling should be defined and maintained at a precise value. Although the master oscillator establishes an accurate and stable phase-lock, the absolute phase of sampling at the adc is influenced by the stability of the various circuits through which the video signal and clock pulses pass before meeting at the sampling gate in the adc. To ensure that the required phase-accuracy of sampling is achieved, a further feedback loop is established around the adc, in which the digitised colour burst is examined; this represents the phase reference for the video chrominance information: see Fig. 7.

At the adc output each cycle of the colour burst will be described by three samples – 1, 2 and 3: see Fig. 8. It is arranged that the correct angle of sampling phase occurs when the digital number corresponding to the magnitude of sample value *2* equals that of sample *3*. These numbers are compared digitally and any differences, which if present indicate a phase error, are accumulated over several burst periods to improve noise performance, and then converted to an analogue voltage using a simple dac. This error voltage is used to vary a voltage-controlled phase-shifter in the reference feed to the master oscillator causing the loop to settle when *2* = *3* and the sampling phase is correct.

It has been found operationally that the incoming NTSC signals may on occasion be subject to transmission and/or encoding errors resulting in an incorrect phase relationship between burst and chroma, giving rise to incorrect hue when the signal

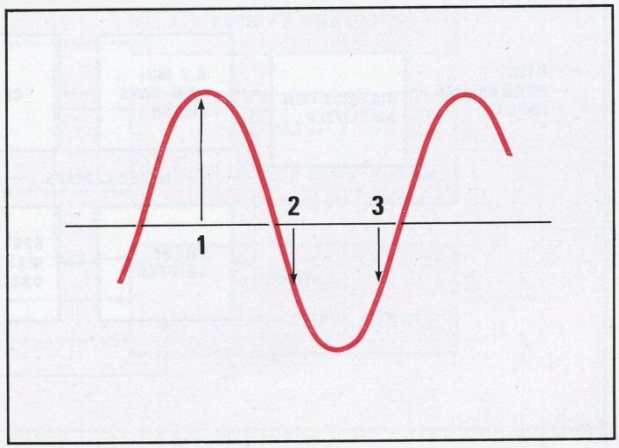

Fig. 8. Each cycle of the colour burst is described by three samples: 1, 2 and 3. It is arranged that the correct angle of sampling phase occurs when the digital number corresponding to the magnitude of sample value 2 equals that of sample value 3.

is decoded in DICE. Since the output from the converter is an encoded PAL or SECAM signal these errors would be difficult to correct later in the transmission chain. Provision is therefore made to allow an operator to over-ride the feedback loop controlling the decoding phase and to correct manually for up to $\pm 40°$ of error, thus ensuring that the converter output is encoded free of hue errors.

525-to-625 Input System

A simplified block diagram of the input signal processing circuitry of DICE operating in the 525-to-625 direction is shown in Fig. 9. The NTSC input signal is terminated in the distribution amplifier (da); this unit has a differential input circuit designed to reject any longitudinal hum that may be present from the picture source. Four isolated outputs are available from the da to drive the various sub-systems and to provide a monitor point for the input programme on the machine monitoring panel. Although DICE is designed to accept a standard level input signal from line, provision is made in the da to correct for level errors of ± 3 dB and chroma/luminance gain errors of 3 dB. The two controls for effecting these corrections are inhibited when the machine is in its calibrated mode but may be activated independently by push buttons located on the control panel.

The main output from the da is filtered to a bandwidth of 4.2 MHz, to remove noise and spurious components outside the NTSC signal spectrum. The filter is of the six-pole elliptic-function type with four stages of group-delay correction designed to give a good stopband attenuation with minimum pulse response error. The output from the filter is then attenuated to bring the signal level to the design value, approximately 400 mV, within the prescribed number of digital levels at the adc. The video is clamped to remove hum and other low frequency errors that may be present, black level being set to -250mV. Since the adc's dynamic range is calibrated to be from -500 mV to $+500$ mV, -250 mV will correspond to the digital number 0100 0000. This clamp also accepts the feedback from the *digital black level stabiliser*.

The digital output signal from the adc, as mentioned

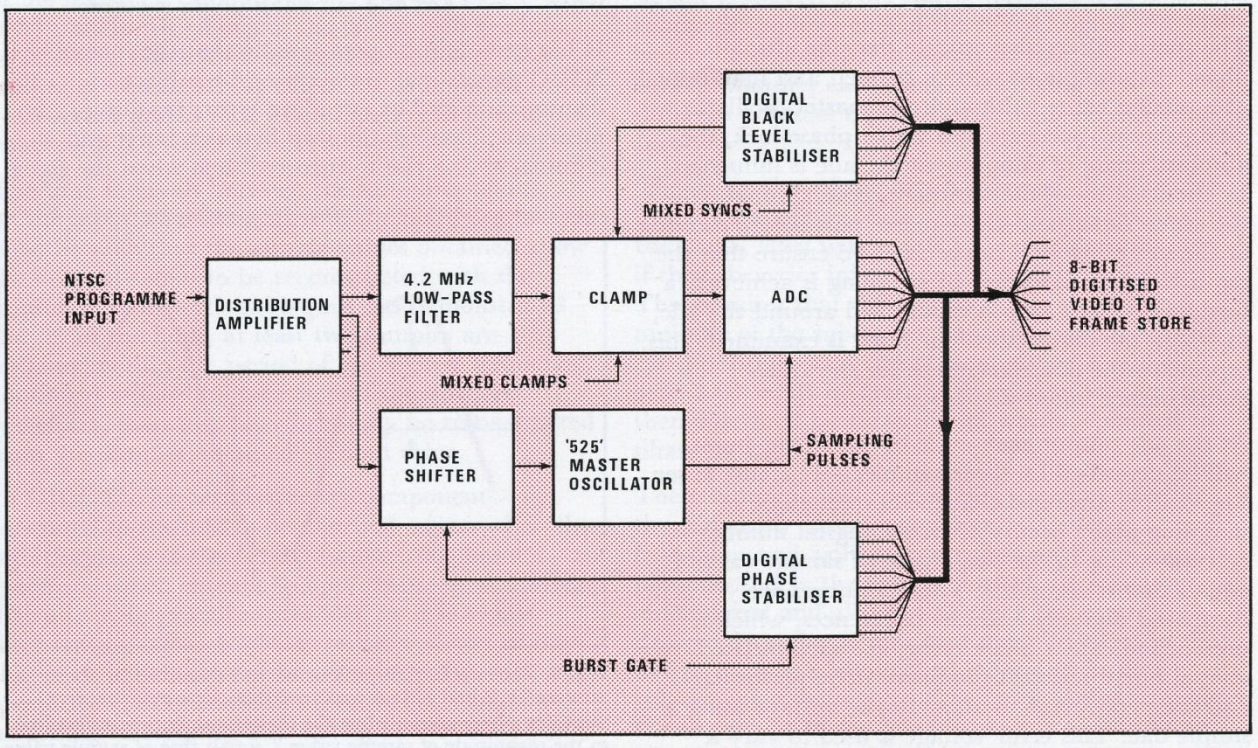

Fig. 9. A simplified block diagram of the input signal processing in DICE when operation is in the 525-to-625 direction.

previously, is eight-bits wide at a data rate of 10.7 Mw/s, and is passed to the frame store. This store has a capacity of one frame of 525-line 60-field per second NTSC encoded video information.

625-to-525 Input System

Since the frame store has a capacity of only one frame of NTSC encoded video information, when DICE is required to operate in the 625-to-525 direction, the incoming 625-line PAL or SECAM signal must be transcoded into a 525-line NTSC format *before* being sent to the store.

The 625-line video input signal from the input distribution amplifier is processed via two different paths, one for luminance and the other for chrominance information. It is in effect decoded down to baseband *RGB*, then matrixed to yield Y, I and Q, the components required to construct the NTSC waveform.

A matrix using the co-efficients

$$E_Y = 0{\cdot}30E_R + 0{\cdot}59E_G + 0{\cdot}11E_B$$

could be used to derive the luminance or Y component of the required NTSC signal: Fig. 10. However, the Y signal resulting from this matrix would contain significant PAL or SECAM subcarrier energy. Standard production decoders, as used to decode the input PAL or SECAM signal in DICE, usually have a simple notch filter in the luminance path to remove or suppress subcarrier components. Although this is adequate to reduce the subjective effect of subcarrier components in highly saturated areas of a viewed picture, it is not sufficient to remove completely the subcarrier sidebands generated on rapid chroma transitions, etc.

As the sampling rate used in the adc ($\simeq$ 10.8 MHz) is asynchronous to residual PAL or SECAM subcarrier components in the derived luminance signal, spurious components would be generated due to 'beat effects' between the two. These would give rise to objectionable patterning on chroma edges, particularly when coded into NTSC by further interaction with the NTSC subcarrier.

To ensure that these possible effects are kept to a minimum, the luminance signal needs to be as clean as possible; to this end it was decided simply to band-limit the composite input video signal using a low-pass filter with a very rapid cut-off, the actual cut-off frequency being a compromise between that

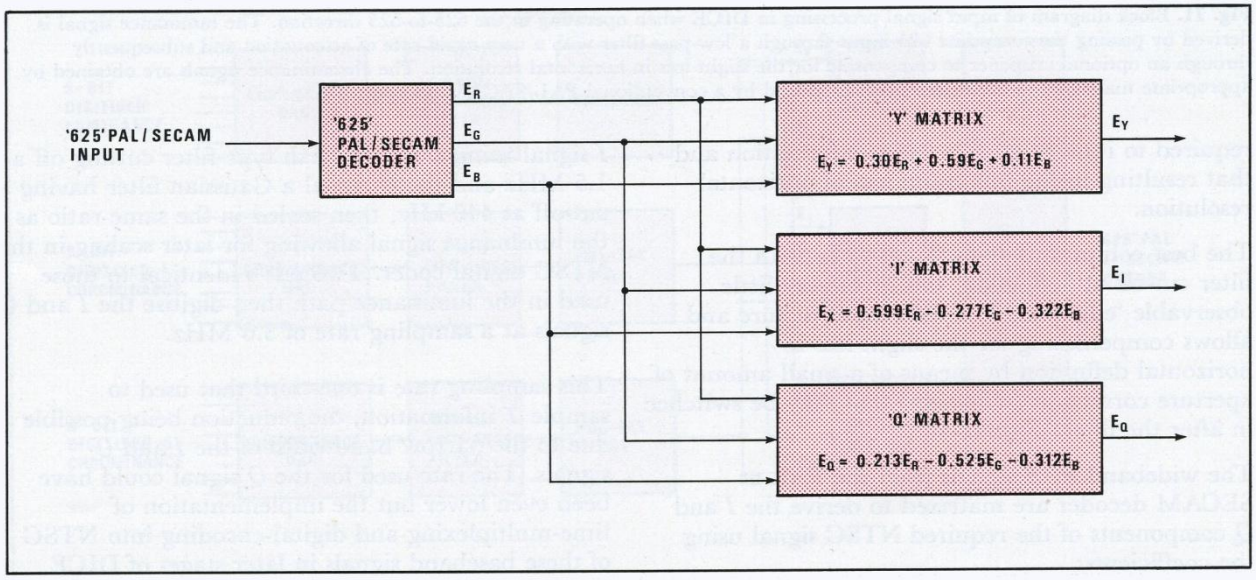

Fig. 10. The Y, I and Q components required to construct the NTSC waveform can all be derived by suitable matrixing of the *RGB* signals produced by decoding the 625-line input signal. However, the Y signal resulting from this matrixing would contain significant PAL or SECAM subcarrier energy and this would cause patterning on the NTSC input signal due to 'beat effects' between the residual 625 subcarrier and the asynchronous clocks used in the 625-to-525 direction, and suitable precautions must be taken.

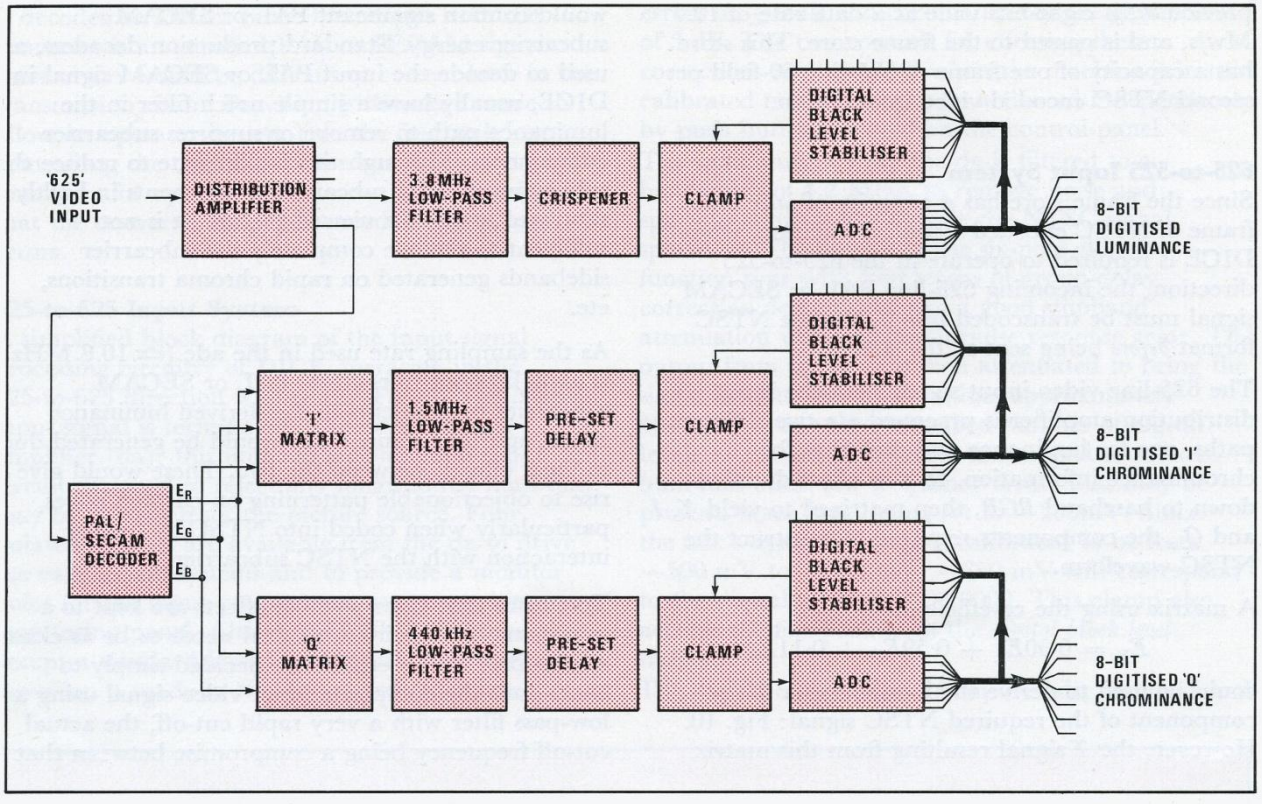

Fig. 11. Block diagram of input signal processing in DICE when operating in the 625-to-525 direction. The luminance signal is derived by passing the composite 625 input through a low-pass filter with a very rapid rate of attenuation and subsequently through an optional crispener to compensate for the slight loss in horizontal resolution. The chrominance signals are obtained by appropriate matrixing of the *RGB* signals produced by a conventional PAL/SECAM decoder.

required to remove all subcarrier information and that resulting in the minimum loss of horizontal resolution.

The best compromise was found to be with the filter cut-off at 3.8 MHz; this gives very little observable 'edge busy' on the output picture and allows compensating for the slight loss in horizontal definition by means of a small amount of aperture correction *(crispening)* that may be switched in after the filter: Fig.11.

The wideband *RGB* signals from the PAL or SECAM decoder are matrixed to derive the I and Q components of the required NTSC signal using the coefficients:

$$E_I = 0.599E_R - 0.277E_G - 0.322E_B$$

$$E_Q = 0.213E_R - 0.525E_G - 0.312E_B$$

The resulting I and Q signals are band-limited, the I signal using a Butterworth type filter cutting off at 1.5 MHz and the Q signal a Gaussian filter having a cut-off at 440 kHz, then scaled in the same ratio as the luminance signal allowing for later scaling in the NTSC digital coder. Two adc's identical to those used in the luminance path then digitise the I and Q signals at a sampling rate of 3.6 MHz.

This sampling rate is one-third that used to sample Y information, the reduction being possible due to the narrow bandwidth of the I and Q signals. The rate used for the Q signal could have been even lower but the implementation of time-multiplexing and digital-encoding into NTSC of these baseband signals in later stages of DICE are eased if the same clock rate is used for the I and Q signals.

The black (or more precisely *zero chrominance*) levels of these signals are set, by further black level

stabilisers associated with each adc, to a digital code 1000 0000 (halfway up the adc transfer curve) giving equal positive and negative headroom for the signals.

Small adjustable delay lines are included in the input feeds to the I and Q adc's, allowing precise adjustment of luminance/chrominance delay to compensate for manufacturing tolerances of the main components in the signal paths, filter, decoder, etc.

The three digital bit streams representing the Y, I and Q components derived from the input PAL or SECAM signal are then interpolated into a 525-line format and subsequently digitally encoded into pseudo-NTSC to feed the frame store.

At the output side of DICE, the signals from the digital circuits of the converter must be translated back into the analogue form and processed where necessary to yield (depending on the direction of operation) the required PAL, SECAM or NTSC output signal.

As an aid to studio 'line-up', provision has been included in the output circuits for colour bars generated within DICE to be substituted for the converted output signal. When the output signal is PAL or SECAM, the colour bars are generated in the output encoder of conventional analogue design; whilst in the NTSC case the colour bars are generated digitally and are switched in at the input to the dac.

525-to-625 Output System

For the 525-to-625 direction of operation, the digital output from the standard converter's logic system comprises three data streams, corresponding to the Y, I and Q components of the input NTSC signal. These signals are converted back into analogue waveforms via three separate dac's, the luminance dac being clocked at 10·8 MHz and the chrominance ones at 3·6 MHz. This is indicated in Fig.12.

The resulting analogue waveforms are filtered to remove out-of-band information and to allow correction of slight errors of luminance and chrominance delay, which may arise due to manufacturing tolerances of filters, etc. Small adjustable delay lines are included in the I and Q paths.

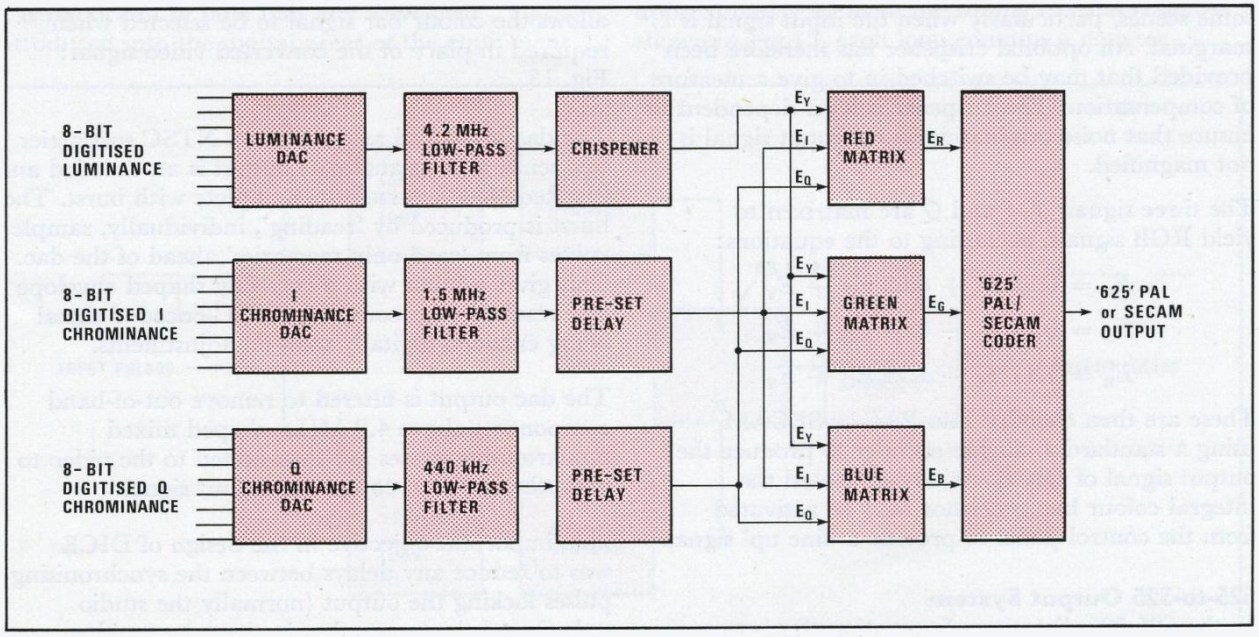

Fig. 12. Output signal processing in DICE in the 525-to-625 direction. The three dac's reconstitute the Y, I and Q baseband signals from which the *RGB* drives for the PAL or SECAM coder are produced by suitable matrices. Out-of-band components in the baseband signals are removed by low-pass filters after the dac's and pre-set delays in the I and Q paths compensate for any errors in luminance/chrominance delay. Any softening of the luminance signal due to processing within DICE may be removed by means of an optional crispener in the luminance path.

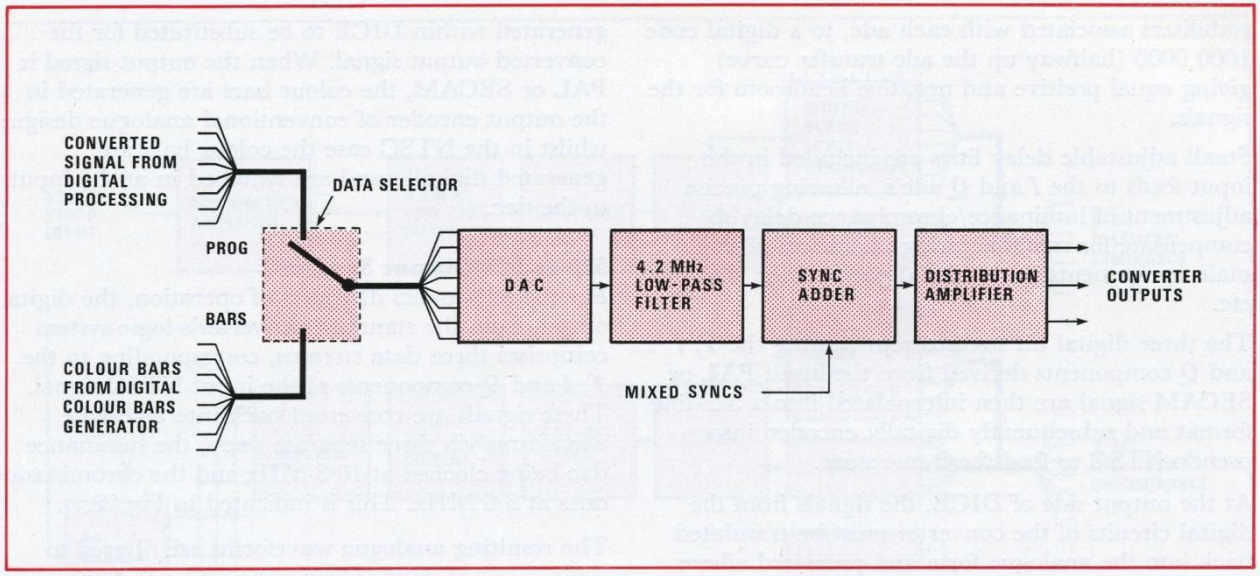

Fig. 13. Output signal processing in DICE in the 625-to-525 direction. A data selector permits selection of either the converted signal from the digital processing of DICE or colour bars generated within the DICE digital colour bar generator. A single dac reconstitutes the non-composite analogue output which is band-limited to 4.2 MHz before the addition of syncs and being 'buffered' to provide three isolated outputs.

Despite the adoption of five-line interpolation in DICE, the output picture may look slightly 'soft' on some scenes, particularly when the input signal is marginal. An optional crispener has therefore been provided that may be switched in to give a measure of compensation. The crispener is level-dependent to ensure that noise originating in the input signal is not magnified.

The three signals Y, I and Q are matrixed to yield RGB signals, according to the equations:

$$E_R = 0.95\ E_I + 0.62\ E_Q + E_Y$$

$$E_G = 0.27\ E_I - 0.64\ E_Q + E_Y$$

$$E_B = 1.11\ E_I + 1.71\ E_Q + E_Y$$

These are then encoded into PAL or SECAM using a standard analogue encoder to produce the output signal of DICE. As already noted the integral colour bar generator may be activated from the control panel to provide a 'line up' signal.

625-to-525 Output System

In the 625-525 direction of operation the output from the standard converter digital system is a single data stream corresponding to a non-composite encoded NTSC output signal, the synchronising pulses being added after conversion to analogue form.

In this mode, digitally generated colour bars are available and a data selector preceding the dac allows the colour bar signal to be selected when required in place of the converted video signal: Fig. 13.

The dac is clocked at three times NTSC subcarrier frequency and its analogue output is an encoded and blanked NTSC waveform, complete with burst. The burst is produced by 'reading', individually, sample values from 'read only memories' ahead of the dac. This gives a burst with a correctly shaped envelope and results in the complete NTSC-encoded signal being encoded digitally with no adjustments.

The dac output is filtered to remove out-of-band components above 4.2 MHz; shaped mixed synchronising pulses are then added to the video to yield the required composite output signal.

One important objective in the design of DICE was to reduce any delays between the synchronising pulses locking the output (normally the studio pulses) and the sync pulses in the output video information to the minimum possible; it was also desirable that any such delays should be the same for both directions of operation in order to allow ready integration of DICE into studio centres.

In the 525-to-625 direction of operation, the only delaying of the sync pulses is due to the PAL or SECAM encoder at the output; such delays amount to approximately 550 ns. This delay is typical of other studio equipment and it was therefore taken as the target value for DICE.

In the 625-to-525 direction, sync pulses are added virtually at the output, but the addition of the colour burst and signal blanking occur ahead of the dac and output filtering; this results in the burst pulse and the blanking being delayed by approximately 2μs before reaching the output. This factor would normally require the sync pulses to be delayed an equal amount in order to obtain a correctly timed output waveform. However in DICE, this problem was resolved by *advancing* the burst gate and blanking.

To obtain this apparent advance, digital phase-locked-loops are used to stabilise the incoming edges of the burst gate and blanking waveforms against the system clocks. The output from these loops consists of edges which are coherent to the system clock pulses with any jitter that may have been present on the input waveform filtered out.

It is important to note that any waveform introduced into the digital areas of the equipment must be coherent with the system clocks in order to prevent the mistiming that would otherwise occur when an input edge jittered across a clock edge: see Fig.14.

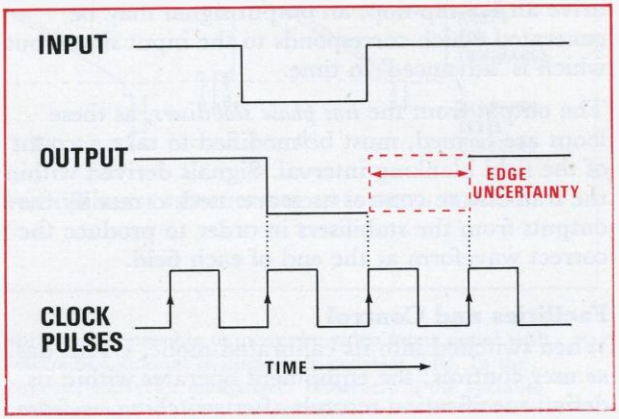

Fig. 14. Any waveform introduced into the digital areas of the equipment must be coherent with the system clocks in order to prevent mistiming which would otherwise occur when an input edge jittered across a clock edge.

The clock rate used in these phase-locked-loops is six times subcarrier frequency, giving an accuracy of approximately 45 ns to the timing of each edge. As shown in Fig.15, each loop contains a counter

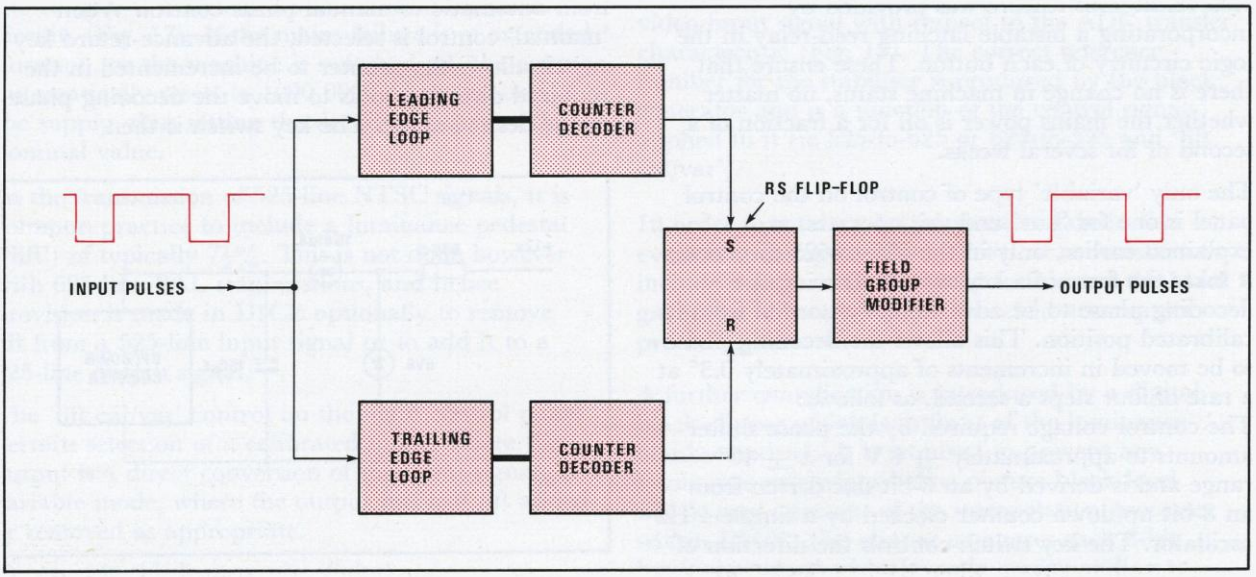

Fig. 15. Digital phase-locked loops are used to stabilise the incoming edges of the burst gate and the blanking waveforms against the system clocks. Two such loops, one locked to the leading edge and the other to the trailing edge of the incoming signal drive associated counters and decoders; in turn these set and re-set an RS flip-flop to reproduce the incoming signal with all edges stabilised. This technique allows an output signal to be produced which is effectively 'advanced' in time with respect to the input.

which can be decoded to give an edge effectively 'advanced' on the incoming edge in steps of 45 ns. By using two loops, one for each edge of the individual waveform, and decoding their outputs to drive an RS flip-flop, an output signal may be generated which corresponds to the input signal but which is 'advanced' in time.

The output from the *line phase stabilisers,* as these loops are termed, must be modified to take account of the field blanking interval. Signals derived within the frame store control area are used to modify the outputs from the stabilisers in order to produce the correct waveform at the end of each field.

Facilities and Control

When switched into its calibrated mode, DICE has *no* user controls; the equipment operates within its design specification seconds after switching on from 'cold'. Nine push buttons are used to select the various functions, whilst a tenth button selects 'remote' operation. Machine status is indicated at all times at the local and remote positions by means of illuminated segments on the buttons.

The logic associated with each control function has been designed to ensure that no change in status can occur during mains transients, surges or failures, or when switching from 'local' to 'remote' operation or vice versa. This feature was provided by incorporating a bistable latching reed-relay in the logic circuitry of each button. These ensure that there is no change in machine status, no matter whether the mains power is off for a fraction of a second or for several weeks.

The only 'variable' type of control on the control panel is one for 'hue' and this operates, as explained earlier, only in the 525-to-625 direction; it takes the form of a key switch allowing the decoding phase to be advanced or retarded from its calibrated position. This allows the decoding phase to be moved in increments of approximately 0.3° at a rate of four steps a second, as follows:
The control voltage required by the phase shifter amounts to approximately $\pm$ 4 V for a $\pm$ 40° range and is derived by an 8-bit dac driven from an 8-bit up/down counter clocked by a simple 4 Hz oscillator. The key switch controls the direction of counting, up or down, when held to the left or right respectively: see Fig. 16.

When decoding along the correct axis, the control voltage is zero volts, hence, when the machine is in

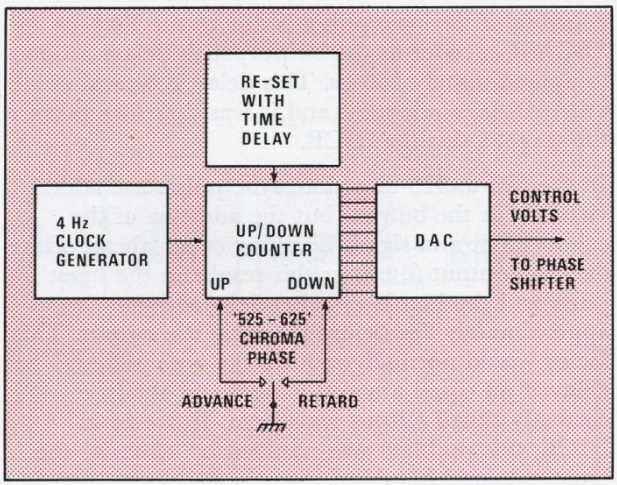

Fig. 16. The only 'variable' control on the DICE control panel is for 'hue' and this is used only in the 525-to-625 direction. It takes the form of a key switch allowing the decoding phase to be advanced or retarded from its calibrated position. The control voltage required by the phase-shifter is derived by an 8-bit dac driven from an 8-bit up/down counter clocked by a simple 4 Hz oscillator. The key switch controls the direction of counting.

its normal 'calibrated' mode, the counter is reset to 1000 0000. This results in 0-V output from the dac and no change in decoding phase when switching from automatic to manual phase control. When 'manual' control is selected, the advance-retard key switch allows the counter to be incremented in the required direction so as to move the decoding phase to the desired angle. The key switch is then

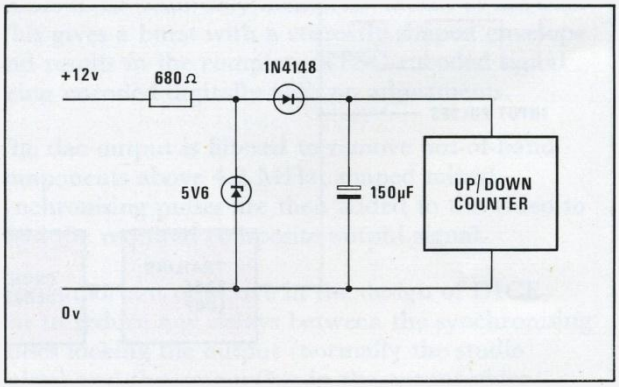

Fig. 17. To protect the counter controlling the phase-shifter from the effects of any mains failures it incorporates a storage capacitor, fed via an isolating diode. To allow this capacitor to be of reasonable size, the counter is based on cmos integrated circuits which present a low quiescent load.

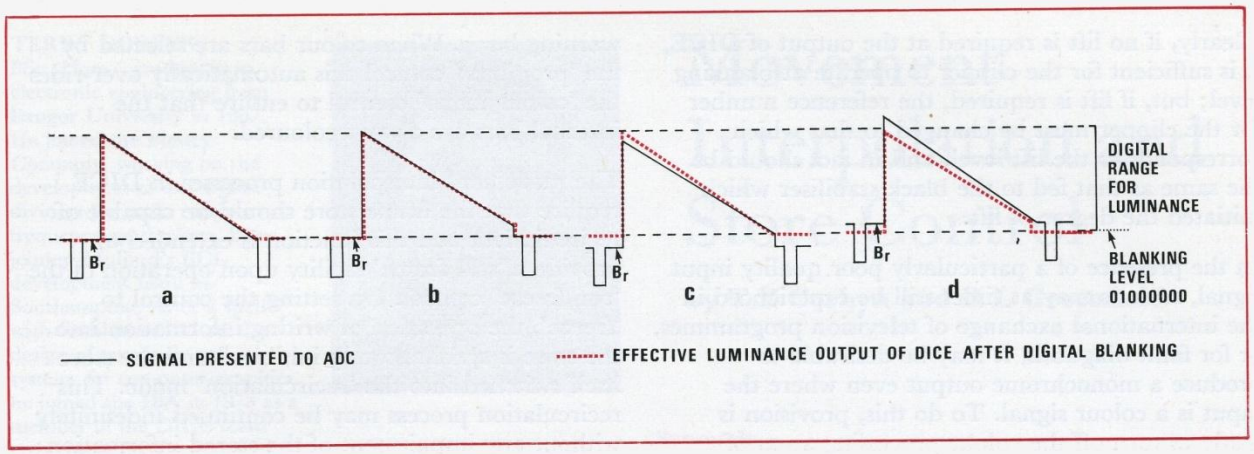

Fig. 18. By varying the reference number to the digital black level stabiliser, it is possible to move the video input signal with respect to the adc transfer characteristic. This adds or removes a predetermined amount of luminance pedestal (lift). In the calibrated mode any signal without lift (a) or with lift (b) passes through the converter unchanged. However, in the variable mode (c) a 525-line input signal with lift present has this removed by referencing the black level stabiliser below blanking level (0100 0000) and at (d) a 625-line input signal without lift has it added by referencing above blanking level. In either (c) or (d) it is necessary to adjust the analogue gain preceding the adc to produce a standard level output.

returned to its 'rest' position and the counter stops, storing the last count.

To protect the contents of the counter against mains failures, it is supplied from a storage capacitor, fed via an isolating diode from the 12 V positive supply. To allow a reasonably-sized capacitor to be used, the counter is designed around cmos integrated circuits (Fig. 17). If the mains failure is of extended duration, or the machine is switched off, the counter automatically resets to 1000 0000 on restoration of the supply, thus setting the decoding phase to the nominal value.

In the transmission of 525-line NTSC signals, it is common practice to include a luminance pedestal ('lift') of typically 7½%. This is not done however with 625-line PAL transmissions, and hence provision is made in DICE optionally to remove lift from a 525-line input signal or to add it to a 525-line output signal.

The 'lift cal/var' control on the main control panel permits selection of a calibrated mode, where the output is a direct conversion of the input signal, or a variable mode, where the output has had lift added or removed as appropriate.

To provide this feature, the digital code corresponding to luminance blanking level at the output of the converter is set to 0100 0000. If the digitised luminance is 'raised' with respect to this code, the effect is to add a fixed amount of pedestal to the output signal. Similarly, 'lowering' the digitised video code removes a fixed amount of pedestal. This facility is illustrated in Fig. 19. In practice, in DICE, the process is implemented by varying the reference number for the digital black level stabiliser; this has the effect of moving the video input signal with respect to the ADC transfer characteristic (Fig. 19). The correct reference number for the stabiliser is produced by the black generator and is a function of the control signals applied to it (ie 525-to-625 or 625-to-525 and 'lift cal/var').

In order to retain a 1:1 conversion gain in the event of a change in the lift level, it is necessary to include a gain correction, and to this end the black generator also controls two switched attenuators preceding the adc.

A further complication is introduced by a digital black clipper which is in front of the luminance blanker and which is required to prevent any luminance excursions below picture black level which may arise out of the interpolation processes within DICE. The clipper compares each 8-bit luminance word with the code corresponding to black level and, in the event of the signal level falling below black level, it supplies a number corresponding to black level.

Clearly, if no lift is required at the output of DICE, it is sufficient for the clipper to operate at blanking level; but, if lift is required, the reference number for the clipper must be changed to that which corresponds to the lift level, and in fact should be the same as that fed to the black stabiliser which initiated the degree of lift.

In the presence of a particularly poor quality input signal, such as may at times still be experienced in the international exchange of television programmes, or for fault diagnosis, it may be desirable to produce a monochrome output even where the input is a colour signal. To do this, provision is made to turn off the colour processing areas of DICE by operation of the 'colour/mono' control. In the case of a monochrome input signal, automatic detection of the lack of colour burst sets the converter to monochrome operation and, if the 'colour/mono' control has inadvertently been left in the 'colour' position, flashes the 'colour off' warning lamp. When colour bars are selected by the 'prog/bars' control this automatically over-rides the 'colour/mono' control to ensure that the internal bars are always coloured.

The movement interpolation processes in DICE require that the frame store should be capable of recirculation, and this function is extended to provide a 'still frame' facility upon operation of the 'run/freeze' control. On setting the control to 'freeze', the operation of writing information into the store is synchronously inhibited and the store itself switched into the 'recirculation' mode. This recirculation process may be continued indefinitely without any impairment of the stored information; DICE continues to provide a still frame output picture corresponding to the contents of the store at the instant of 'freezing'. The main purpose of this 'run/freeze' control is to aid the diagnosis of faults should they occur, eg a fault before or in the writing of the frame stores can be frozen.

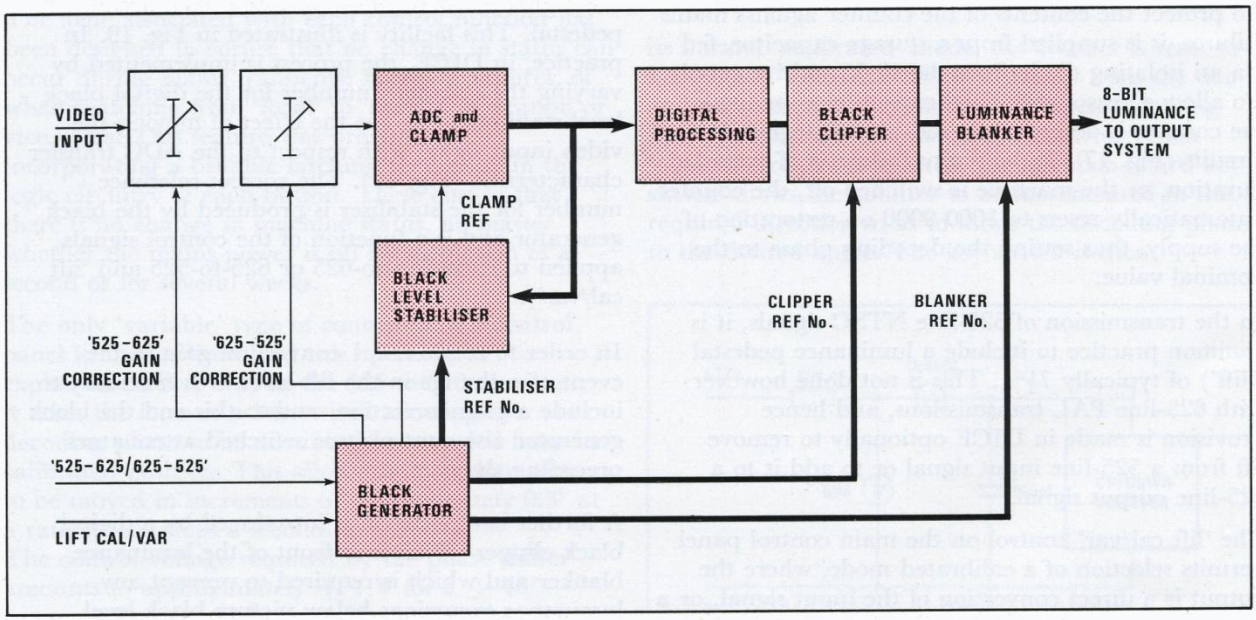

Fig. 19. Block diagram of the system for adding or removing lift. When 'lift var' is selected the black generator changes the reference number to the black stabiliser, thus adding or removing lift according to the direction of conversion. The amount of variation in lift level is governed by pre-set switches on the black generator and normally set to add or remove 7½%. Gain correction commands are also produced by the black generator to control switched attenuators preceding the adc which restore the conversion gain to 1:1 for a 7½% change in lift. When lift is added to a 525-line output signal the black clipper, which prevents any excursions below picture black which may arise in the digital processing of DICE, is referenced to the same level as the black stabiliser instead of to that of the blanker, selection of the reference level being carried out by the black generator.

Movement Interpolation and Store Control

by T E Corbyn and R L Greenfield

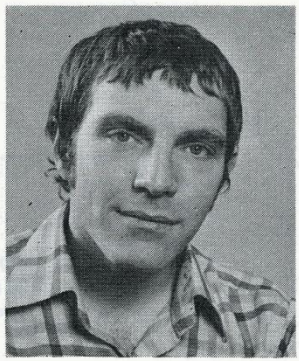

TERRY CORBYN, BSc (Hons), graduated in electronic engineering from Bangor University in 1967. He joined the Plessey Company, working on the development of uhf digital dividers for radio receiver frequency-synthesisers, later joining Mullard's ECL development team at Southampton. After a spell with Stabletron Ltd on the design of terminals and systems for computer-graphics, he joined the IBA in 1973 as a member of the DICE team responsible for the design of the store control system.

RAYMOND GREENFIELD, MIEE, served a student apprenticeship with the Marconi Company, subsequently working on the design of television studio equipment. He later joined Advance Electronics where he was responsible for the design of pulse generators and storage oscilloscopes. Since joining the IBA in 1970 his work has centred around digital television with particular application to DICE.

Synopsis

Digital storage offers significant advantages in the field of standards conversion, allowing signals to be reconstituted without impairment after any required period of time delay, or in a different sequence, or at a different rate.

The section discusses the problem of movement reproduction when converting between two television systems with different field rates. By using the information from two consecutive input fields it is shown how the visible result is improved. In order to provide these two input fields some form of field storage is required. The availability of low cost mosfet shift registers provides means whereby this may be achieved. The operation of these devices is described together with their application to the field store.

The organisation of the two field stores allows reading and writing of the data to take place in a field store with an overlap of the two operations occurring for only a short time. It is arranged that when overlap does occur no reading takes place from this particular field store and that the other field store then provides 100% of the output signal.

The control of the various store operations is described together with the method of changing the number of lines per field in either direction of conversion.

IMAGINE that a converter is operating in the 525-to-625 direction, and that the input picture consists of an object moving across the screen at a rate of 1cm for every input field. After 1/10 second (ie six input fields) the object will, therefore, have moved 6cm. Clearly, the same situation must prevail at the output. However, the output rate is 50 fields per second; so in the same time period as six input fields, five output fields would have been delivered. If the output picture were merely a time stretched input field, then the observed velocity of the output would be 5cm for every tenth of a second. Not only would this be true, but also all converted television programmes would last 20% longer than the unconverted programme!

Since there are six input fields for every five output fields, a simple way of obtaining an average output movement equal to the input movement is to omit every sixth input field. Fig. 1 shows the distance/time graph of an object moving 1cm per input field. For four output fields the velocity is slowed to 5/6cm per input field, and for the fifth output field when the omission occurs, the output velocity is suddenly doubled. On average, therefore, the output distance/time graph follows that of the input.

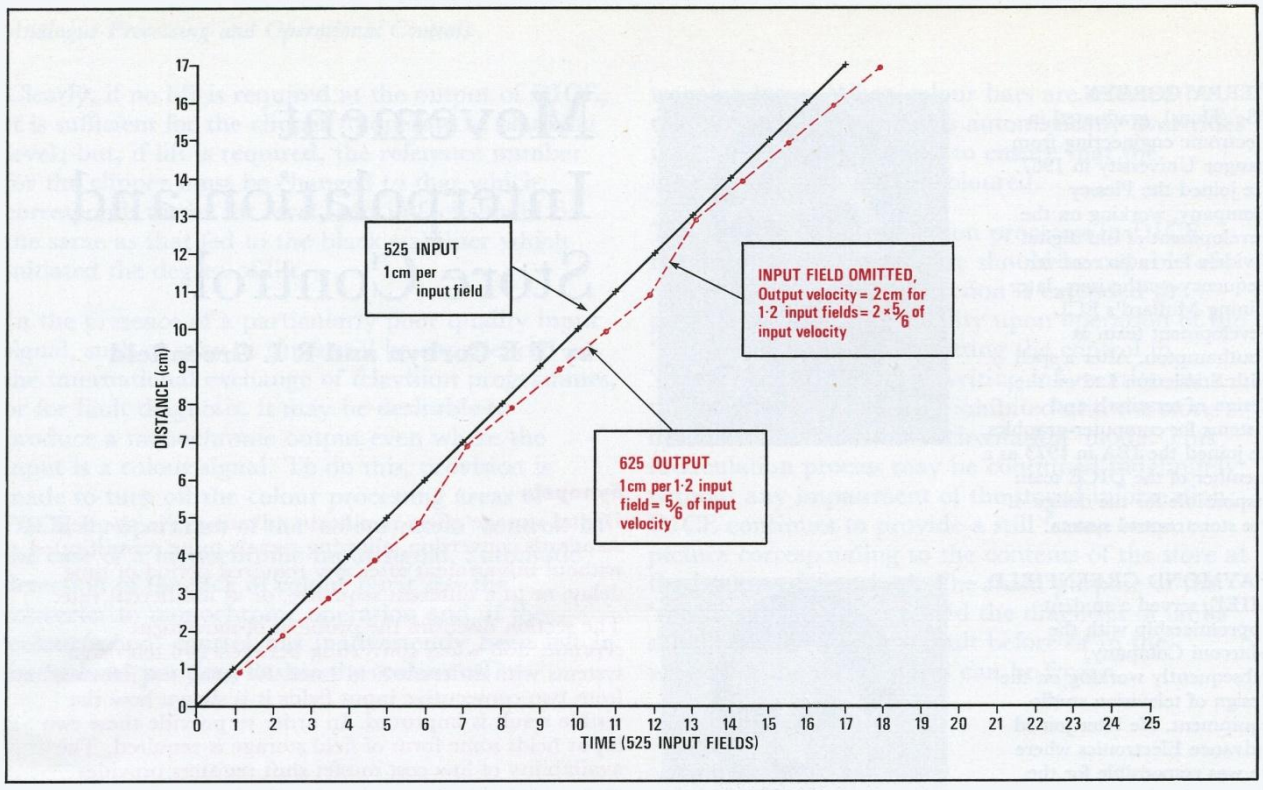

Fig. 1. Removal of input fields so that average output velocity equals input velocity.

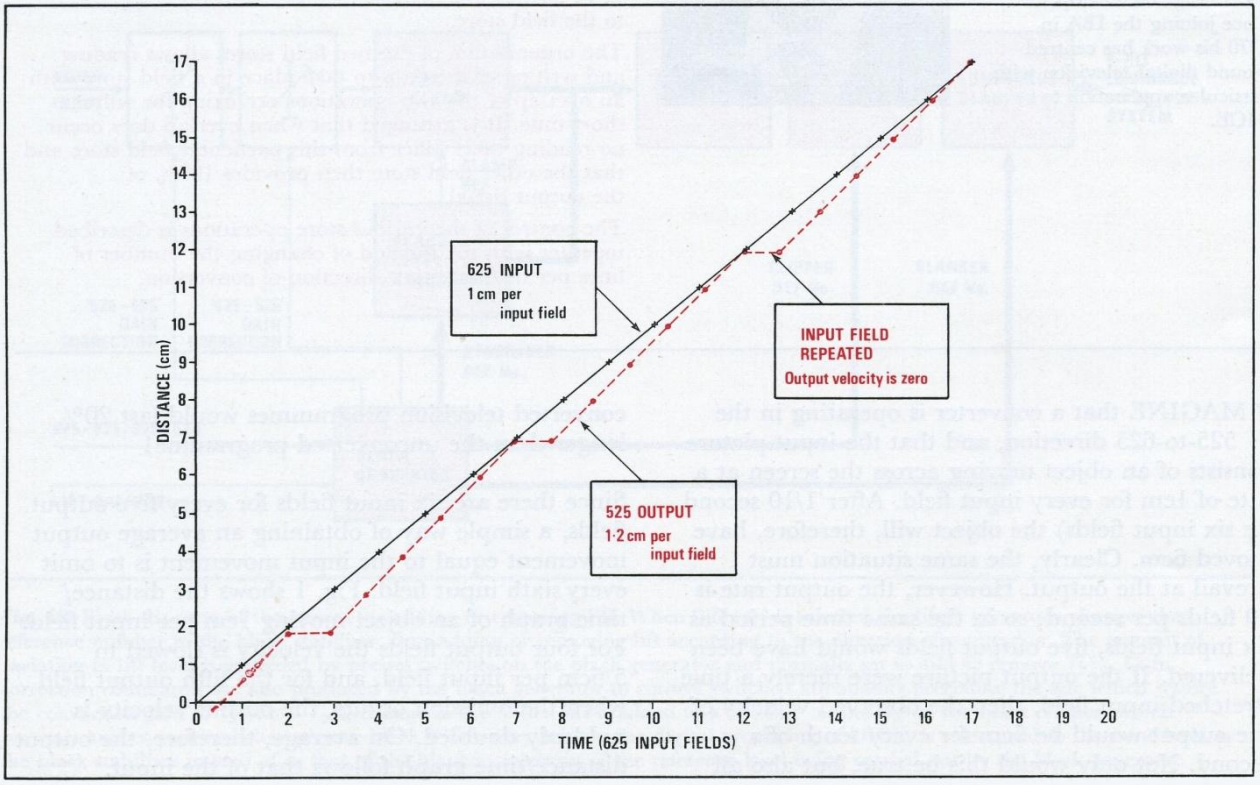

Fig. 2. 625-to-525 direction, without interpolation. Every fifth input field is repeated at the output. Average output velocity is equal to input velocity.

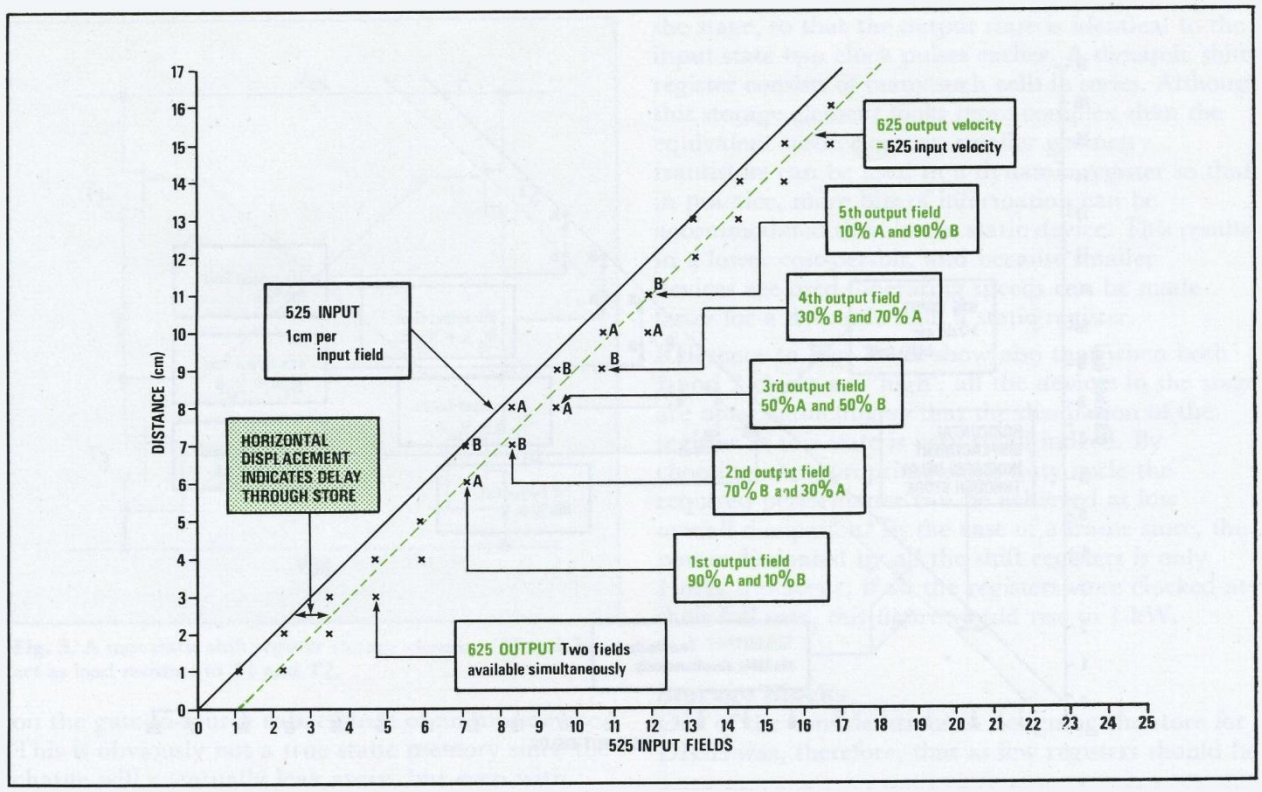

Fig. 3. 525-to-625 direction. Use of movement interpolation to mix the outputs of the two stores in varying proportions so that output velocity equals input velocity.

However, the visual effect at the output is most objectionable, since the omission of input fields appears as distinct 'jumps' in the output picture. In the 625-to-525 direction a similar stratagem leads to every sixth output field being repeated: see Fig. 2.

A better approach to the problem is to combine the outputs of two field stores in various proportions. The proportions are varied according to the relative timing of input and output fields. Fig. 3 shows this operation in the 525-to-625 direction. Note that the two simultaneously available fields are mixed using the coefficients 90%, 70%, 50%, 30%, 10%.

The corresponding situation for the 625-to-525 direction is shown in Fig. 4. Here, the motion of the image at the input is assumed to be 1.2 cm per input field (since this shows the required coefficients more clearly). In this case six coefficients are used: 12/12, 9/12, 7/12, 5/12, 3/12 and 0/12.

The results obtained from movement interpolation are not perfect, especially at certain speeds of movement. Nevertheless, the observed picture is better in terms of movement judder than the picture from a conventional telecine.

Main Store Fundamentals

In order to obtain the time redistribution necessary for standards conversion some form of storage is required to store each incoming field, and then, at some later time, deliver the same field at a rate related to the output standard. Further, the requirements of movement interpolation, as noted above, dictate that at least two fields (ie one complete frame) of the incoming picture should be stored so that they can become simultaneously available at the output of the main store. The main store in the DICE converter consists of two identical field stores, each capable of storing an active field (ie that part of the field containing video information) of a 525-line NTSC picture. Although one of the television standards handled by DICE is 625 lines, it is only necessary to store a 525-line frame, since, in the 525-to-625 direction this is the

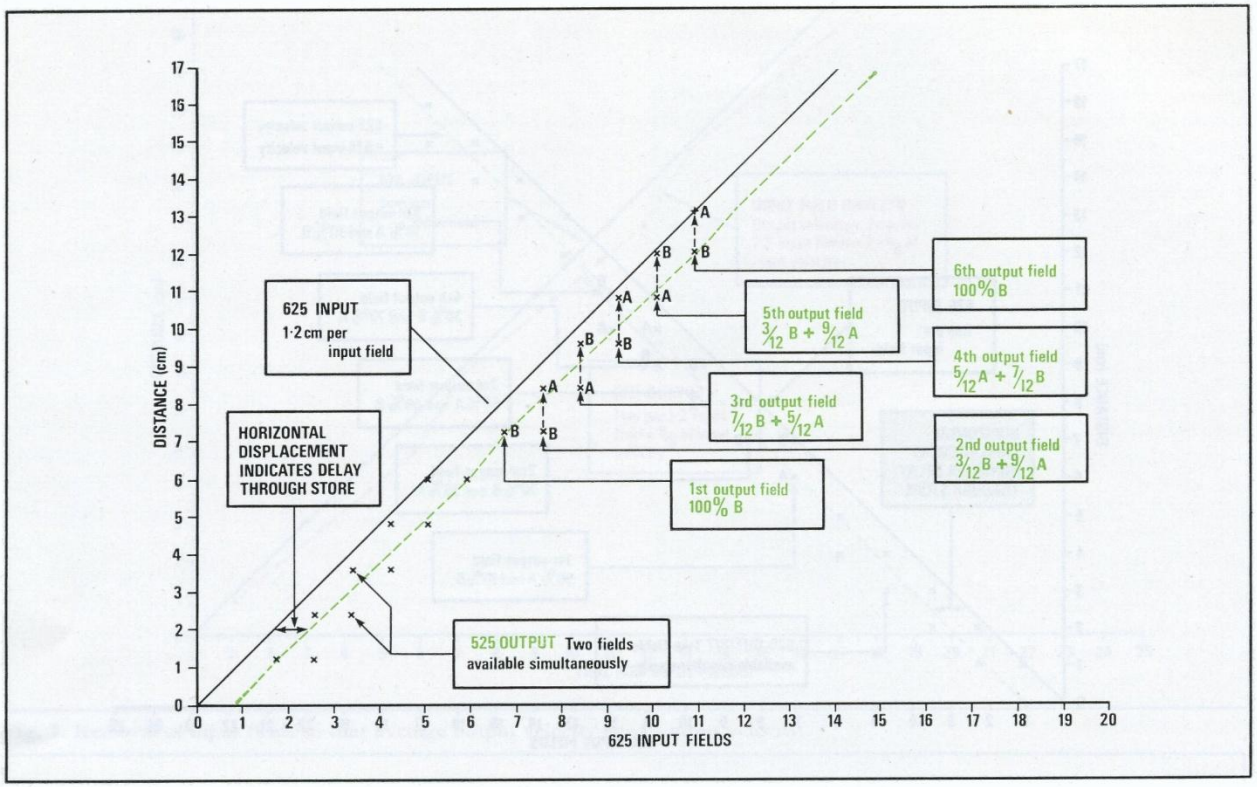

Fig. 4. 625-to-525 direction. Use of interpolation to make output velocity more uniform. Outputs of stores are mixed in various proportions.

total amount of information available at the input, and in the 625-to-525 direction it is the amount of information required at the output.

Storage Elements

Because of their low cost and ruggedness the storage elements used are mos shift registers. These are serial-in, serial-out devices, which make them very convenient for television applications where data is presented to the stores sequentially and is also required at the output sequentially. When data is stored it is shifted into the registers. Having filled the registers with data, the shifting operation can be stopped and data can then be held for a time. When the output is required, data can be shifted out of the register at any required rate. The write-in, storage and read-out properties make the shift register an ideal device for standards converters. A 525-line colour signal sampled at about 10.7 MHz, using an 8-bit word for each sample, requires about 2.4 megabits of storage for two fields.

The best known binary storage element is the bistable circuit. When implemented in *p*-channel mos it has the form shown in Fig. 5.

A pair of cross-coupled inverters, T1 and T2, with mos load resistors, T3 and T4, will remain in one of two possible states with one 'on' and the other 'off' for as long as power is applied. This is a *static memory*. It has the disadvantage that the cell size (ie the size of the bistable element when fabricated on an integrated circuit) is so large that comparatively few such cells can be accommodated on a normal sized silicon chip. This limits the amount of data which can be stored in such a register. Additionally, there is a speed limitation when switching the bistable. Unless device fabrication technologies other than the standard *p*-channel process are used, the shift rate of such a register is unlikely to be above 2 MHz.

Since the input impedance of an mos transistor is extremely high (about 10^{12} ohms) a short term memory can be realised simply by placing a charge

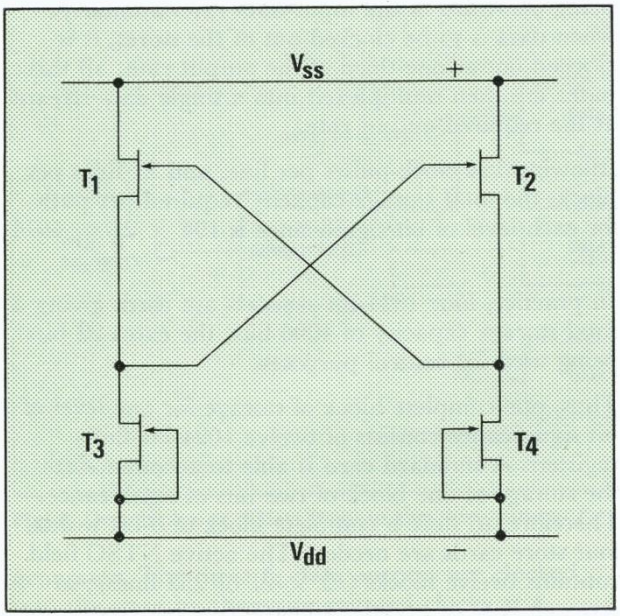

Fig. 5. A mos static shift register storage element. T3 and T4 act as load resistors to T1 and T2.

on the gate-to-source capacitance of an mos device. This is obviously not a true static memory since the charge will eventually leak away, but even with standard devices, information can be retained for several tens of milliseconds.

This is called a *dynamic memory*, and the principle is shown in Fig. 6. Note that the load fet can be disconnected from the negative rail without affecting the charge stored on the gate. However, there are still leakage paths; unless the charge is replenished the information will be lost. This point must be taken into account when considering the field stores.

In practice, the circuit diagram of a single storage element is much more complicated than shown in Fig. 5. Fig. 7 shows the diagram of a practical circuit arrangement for a single stage of an mos dynamic shift register.

There are six mosfets: T1 and T4 are inverters; T2 and T5 act as load resistors; and T3 and T6 are coupling switches. The clock consists of negative-going pulses which are applied alternately to the 1 and 2 inputs of the register. When 1 goes negative the input state is inverted by T1 and T2 and is transferred by T3 to the gate capacitance of T4. Similarly, when 2 goes negative the pulses are inverted and transferred via T6 to the next stage. The data is thus inverted twice in passing through the stage, so that the output state is identical to the input state two clock pulses earlier. A dynamic shift register consists of many such cells in series. Although this storage element looks more complex than the equivalent static element, smaller geometry transistors can be used in a dynamic register so that, in practice, more bits of information can be accommodated than in the static device. This results in a lower cost-per-bit, and because smaller devices are used, operating speeds can be made faster for a dynamic than a static register.

Reference to Fig. 7 will show also that when both 1 and 2 clocks are 'high', all the devices in the stage are non-conducting so that the dissipation of the register in this state is very small indeed. By choosing an appropriate clock duty cycle the required performance can be achieved at low overall dissipation. In the case of a frame store, the power dissipated by all the shift registers is only 1 mW quiescent; if all the registers were clocked at their full rate, this figure would rise to 1 kW.

Storage Blocks

One of the considerations in designing the store for DICE was, therefore, that as few registers should be

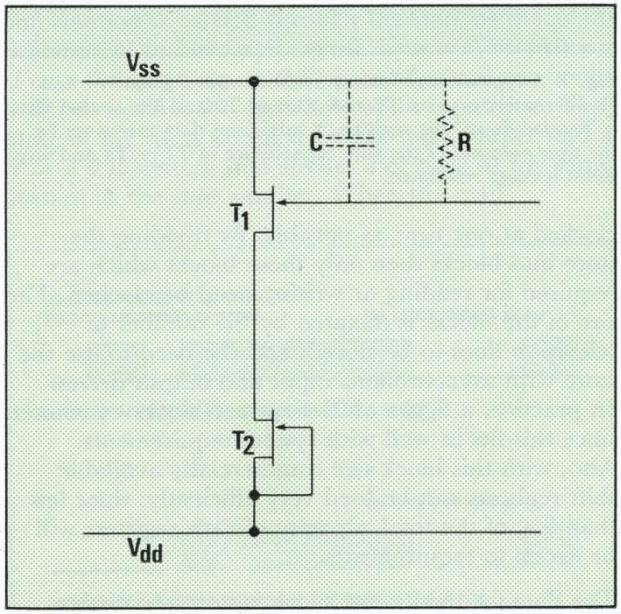

Fig. 6. Temporary storage of charge on the gate to source capacitance, *C*, of mos transistor T1. *R* represents the leakage path of the capacitor. Since much of the leakage takes place through reverse-biased junctions, the value of *R* falls rapidly with increasing temperature.

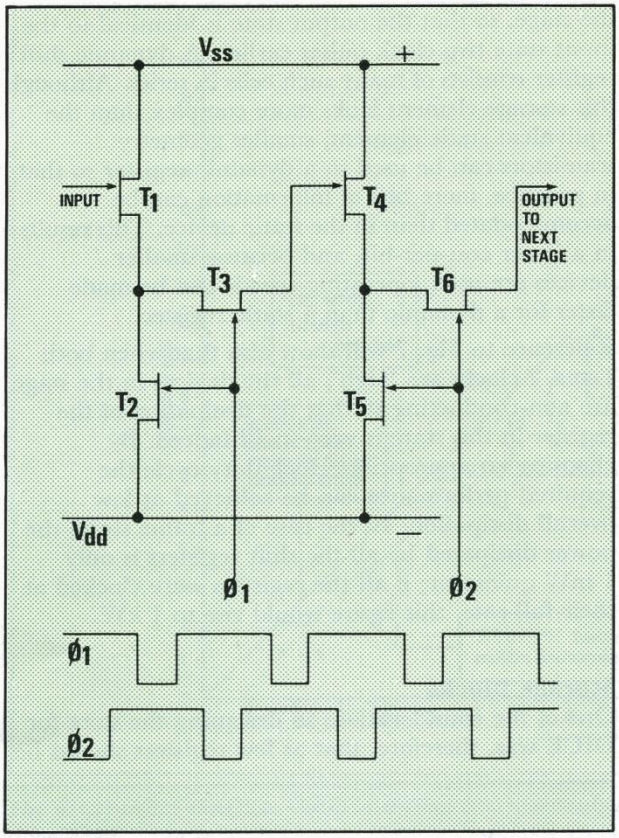

Fig. 7. A typical stage from a dynamic shift register. $\phi 1$ and $\phi 2$ alternately turn on T2 and T3, and T5 and T6, so that data is 'handed down' the register. A minimum data rate is set by the time constant of the gate capacitance of T1 or T4 and any shunt leakage resistance.

clocked at one time as possible. By dividing the store into blocks then only those blocks which are required for reading or writing need be clocked. The size of the blocks is dictated by the number of television lines to be stored, and the need to use the store with asynchronous input and output drives. In practice, a figure of 21 lines per block was chosen since this fits in well with system requirements. Also, with this block size commercially available shift registers can be used very efficiently, since less than 2% of the total capacity of such registers will be excess to requirements.

The date for the 21-line block is stored in twelve 1024-bit registers. Since the input data rate (10.7 MHz) is too high to be handled directly by the registers, the data stream is first demultiplexed into three parallel paths, the data rate on each path being one-third of the input data rate. Similarly, when data is to be clocked out of the stores, it is clocked out at one-third of the output rate; all three paths are then multiplexed into a single data stream at the required output rate.

Allowing 194 clock-pulses per line, and 21 lines per block, the total storage capacity required per path for each level of bit significance is $194 \times 21 = 4074$ bits.

In practice, four 1024-bit registers are used, giving a total storage capacity of 4096 bits, the extra 22 bits being used for control purposes.

Thus, one complete block of storage for one level of bit significance consists of twelve 1024-bit shift-registers. A standard card is able to accommodate the storage of one block of two bits of significance. Therefore, to store a complete block of 8-bit words, four store cards are needed. The active NTSC field contains twelve blocks (12×21 or 252 lines); so for each field of storage, 48 cards are required.

Operation of Main Store

As the picture information arrives field by field, in an odd/even sequence, each sequence of these fields is written into alternate stores. Store selection for 'writing' is under the command of the 'write generator' in store control. Either store may hold an odd or an even field, depending on the start of the storage sequence. When the stores are read, however, it is necessary to know which store holds which type of field, since this will determine the spatial relationships in the data from the two stores. A control signal called *store sense* is provided to perform this function. The store sense is divided into two separate lines; *store sense right* and *store sense left*. In normal operation these lines are complementary, but when a situation occurs where the odd/even sequence is interrupted, ie during a 'cut' from one video source to another, then the sense lines may temporarily assume the same polarity.

Because the timings of write and read signals are completely asynchronous, it is possible for conditions to arise under which reading and writing could be required to take place in one store block simultaneously. Under these conditions the 'read' operation is inhibited to allow writing to proceed continuously.

Data information is clocked into the demultiplexers by the $3f_{sc}$ write-counted clocks from the write

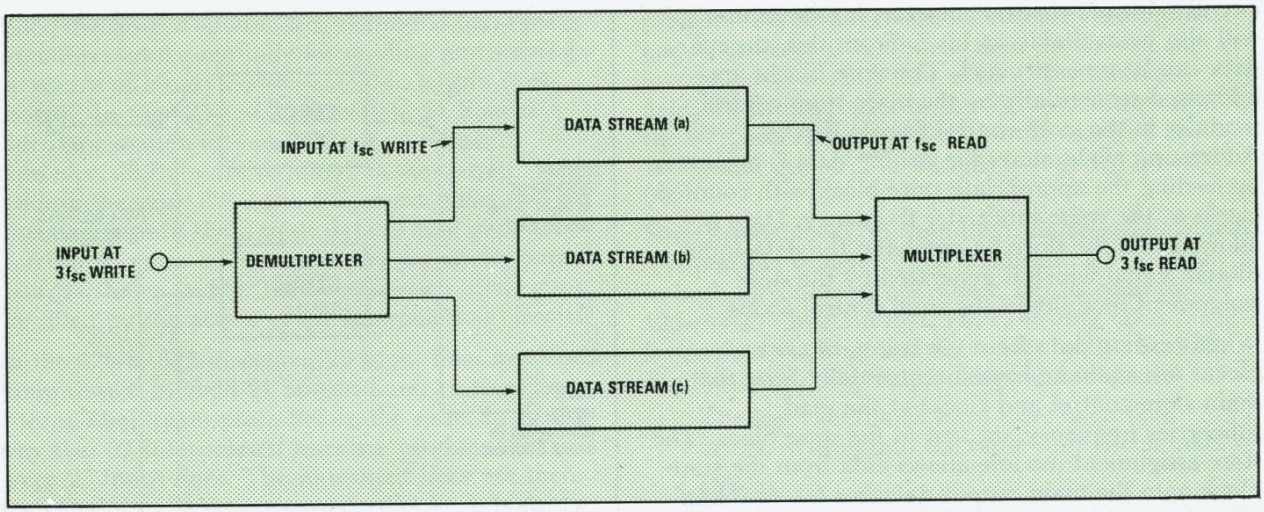

Fig. 8. Demultiplexing of data input to give three paths at f_{sc} write, and recombination of data paths to give $3f_{sc}$ read.

generator. The three f_{sc} data outputs are clocked into the store by the write-counted clocks at f_{sc}, where it remains until it is required to be read. The time which elapses between writing and reading data could theoretically be as long as an output field period, which may be 20 mS in the 525-to-625 direction. This period is too long for the registers in the main store to retain data reliably, especially at elevated temperatures, so a low frequency 'refresh clock', at about 1 kHz, is then applied to those registers which are neither reading nor writing. This ensures that data information will not be lost due to the leakage within the mos devices. Unfortunately, it does mean that when data information is required from the store, its location within the shift register is no longer known. To overcome this problem, a *data start marker* is written into the store in the location immediately ahead of the data. The start marker consists of a logic *1* which is written into data stream *a* (Fig. 8) of the most significant bit (msb) of stored data. It is written by the demultiplexer, under the command of the write generator. At the output of the store the read generators are able to locate the start marker and hence identify the start of video data.

When writing into the store, it is important that the start marker be written unambiguously in data stream *a*. There will still be data in the store after the previous operation and it is quite possible that there will be several *1*s written around the location of the start marker. To obviate this difficulty, the first 23 locations in the block are filled with an *0* at the beginning of every block. Then the start marker is written. When this has been done, the first *1* located in bit *a* of the store block *must* be the start marker: Fig. 9.

For the correct digital decoding of the chrominance information it is important to know the phase of the chrominance on the line under consideration. This information is carried through the store as the *subcarrier phase ident* (spi) signal. This will be either a *1* or a *0* depending on the phase of subcarrier, and will be constant for the whole of an input line. The state of the spi is sampled by the start marker and a *1* or a *0* is written into stream *a* of bit 2 at the same instant as the start marker. When the read

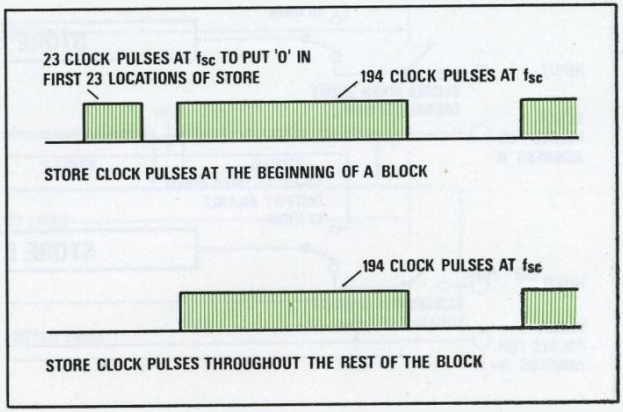

Fig. 9. Write f_{sc} clocks to main store, showing additional 23 clock pulses to clear first 23 locations so that the start marker is unambiguous.

circuitry locates the start marker, it follows that spi will also be available on bit 2. From this, the spi data can be reconstructed. The store blocks are addressed sequentially by the main store control modules in the store control frame: Fig. 10. Addressing of the stores for writing data is achieved by feeding the appropriate store block with counted clocks at the write subcarrier frequency. Data is allowed into the store by removing the data input disable signal which is generated by the store assignment modules in the store control frame.

When reading data from the stores, the store blocks are again addressed sequentially from the main store control and clocks at the read subcarrier frequency are sent to the store block. The store assignment module allows data from the store block onto the main output data bus by 'enabling' tri-state gates at the output of the store. Also, to ensure that the data read-out is non-destructive, the input to the store block is connected to its output so that recirculation of data takes place. This is necessary for movement interpolation, since the contents of a store are often required for two consecutive output fields.

A simplified diagram of a store block is shown in Fig. 11. Note that all clock pulses (write, read and refresh) are sent to a store block down one wire, the combination of the clocks taking place in the main store control modules.

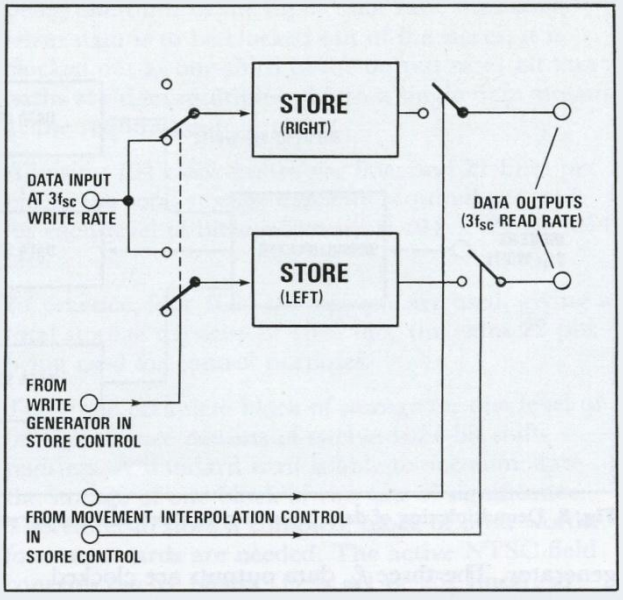

Fig. 10. Main store. The writing carries on continuously and reading may be inhibited by movement interpolation control.

Store Control

The function of 'store control' is primarily to look after the operation of the main store; additionally, however, it delivers various other control signals to other parts of the converter.

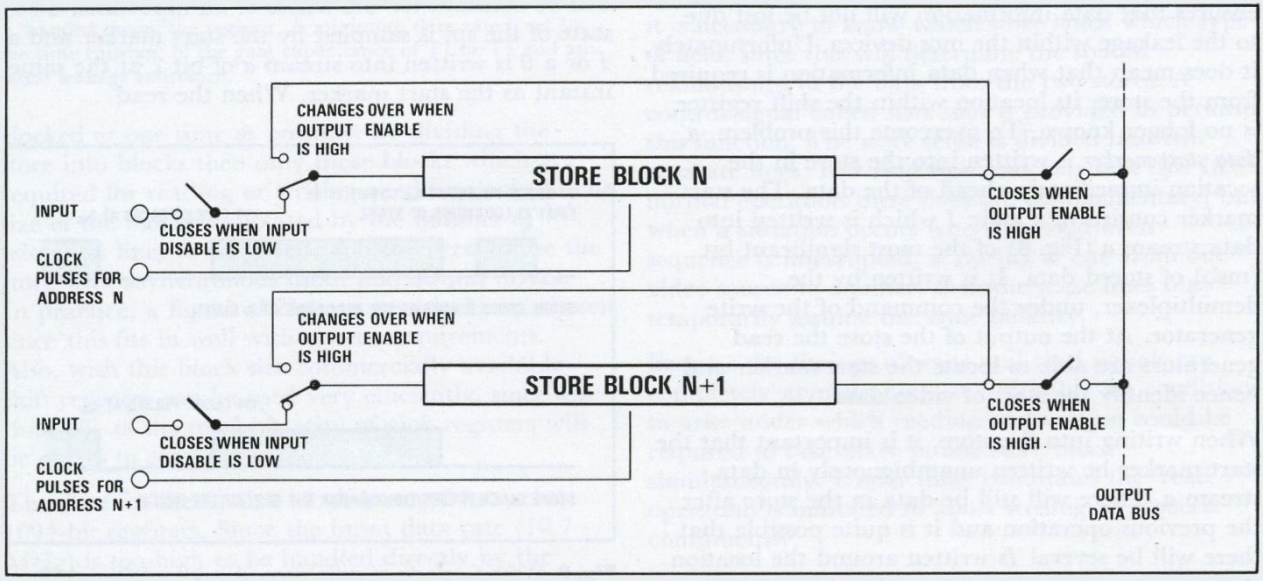

Fig. 11. Organisation of store blocks. When blocks are read, output enable is high and data recirculates. When data is written, switch at input closes and data is allowed into store. Reading and writing do not occur simultaneously.

Since the store is subdivided, the individual blocks are addressed sequentially for reading and writing by means of binary codes generated within store control. In the 525-to-625 direction, 21 lines of video are written into each block of storage. However, 25 lines are required at the output. When reading the store the extra 4 lines are generated by inserting gaps in the output data. These gaps are distributed more or less uniformly across the output data, as shown in Fig. 12. Subsequently these gaps are filled by the action of the interpolation circuits.

In the 625-to-525 direction, 25 input lines must be 'compressed' to yield 21 output lines. In this case, interpolation takes place before the store, in such a way that 21 interpolated lines are produced for every 25 input lines. The unwanted lines are then discarded during the write operation by inhibiting the write clocks when the unwanted lines occur. As explained earlier, because the machine is operating asynchronously, there is the possibility of simultaneous reading and writing taking place in the same store block; however, reading can be inhibited so that writing may be an uninterrupted process. It is arranged that whenever reading of a particular store is inhibited, then the movement interpolation circuits take none of the inhibited field. To facilitate this in the 525-to-625 direction, the 90% interpolation coefficient is rounded up to 100% and the 10% coefficient is rounded down to 0%.

Since a store block is the minimum addressable element, it is possible to deduce the conditions for 'fouling' in both directions of operation. Let us consider the 525-to-625 direction first. Here, writing takes place faster than reading, and there will be two limiting conditions under which fouling cannot occur:

(1) The last store block must have been read before writing into the same store block takes place.

(2) The first store block cannot be read until it has been written into.

Condition (1) summarises the 'overtaking' of reading by the writing operation and (2) shows the writing operation receding away from the reading.

In the 625-to-525 direction, reading is faster than

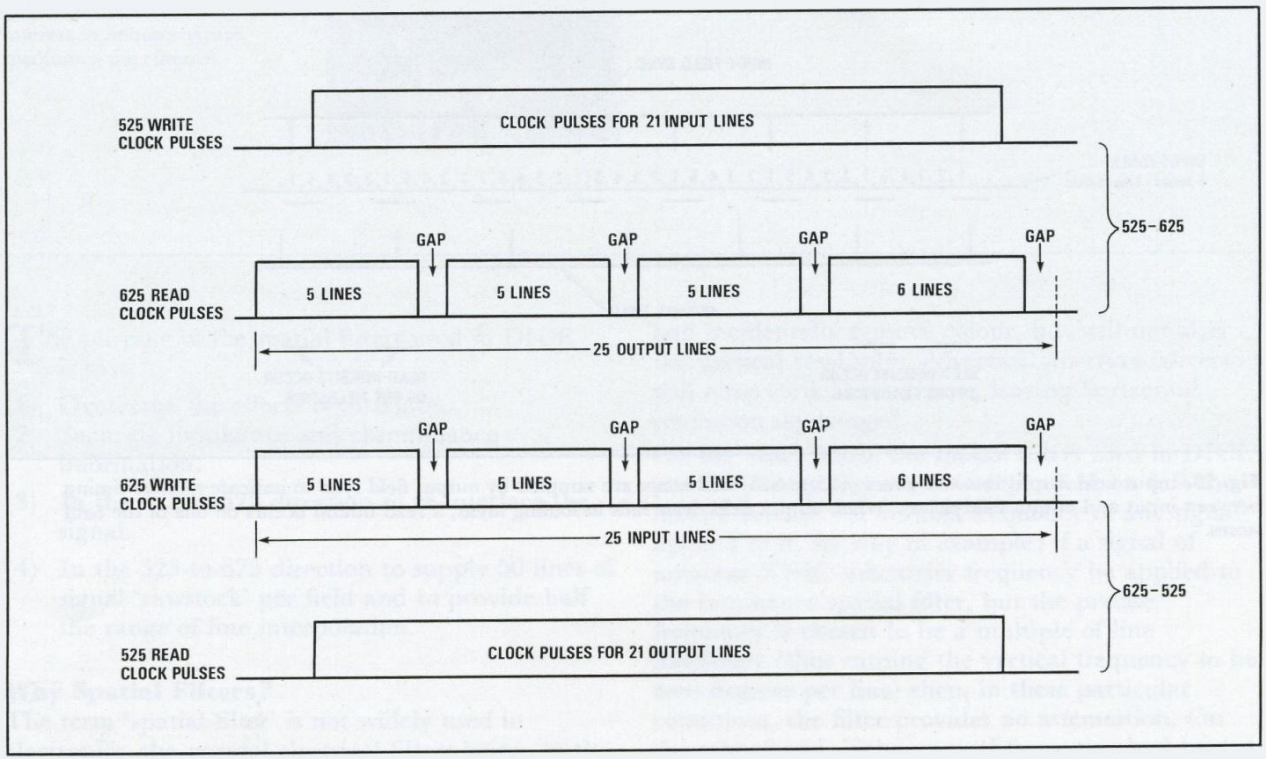

Fig. 12. Read and write clock waveforms showing addition of gap sequence from interpolation control for addition of lines in 525-to-625 direction and removal of lines in 625-to-525 direction.

writing, and therefore the limiting conditions are the reverse:

(1) The last store block must have been written into before reading takes place.

(2) Writing cannot occur in the first store block until it has been read.

If the timing diagrams are drawn, then it will be found that there is a region about 100 lines wide, positioned about the write field sync, which corresponds to the fouling zone. If the read field sync falls within the fouling zone then the reading from one or other of the main stores must be inhibited. The other combinations of coefficients are generated by detecting the dynamic phase relationship between reading and writing in the following way:

The input field period is divided into a number of equal (or as nearly equal as circumstances permit) 'zones'. The number of zones will vary according to the direction of conversion of DICE, five zones being used in the 525-to-625 direction and six zones in the 625-to-525 direction.

A counter/decoder generates the time sequential zones and the read field sync samples the state of the decoder. Thus, the phase relationship for that output field is determined. Fortunately, the fouling area is just less than two zones wide; by phase shifting the zones with respect to the input field sync so that two whole zones occupy the fouling area, the 'read inhibit' signals can also be generated. For instance, in Fig.13, it may be seen that when the read field sync falls in zones one or five, then read inhibiting must take place, and the output coefficients are the 0% or 100%.

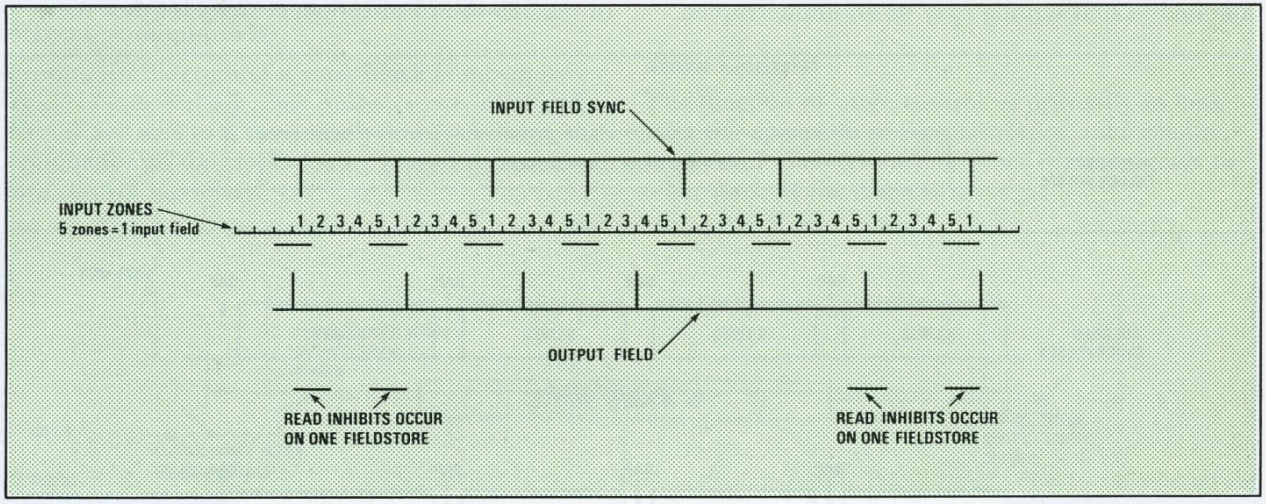

Fig. 13. Input field is split into five zones (525-to-625) and zones are sampled by output field sync to indicate relative timing between input and output field drives. When output field drive falls in fouling areas, a read inhibit occurs on one of the field stores.

JOHN L E BALDWIN, BSc, Head of Video and Colour Section of the IBA's Experimental and Development Department. A biographical note appears on page 3.

Spatial Filters

by J L E Baldwin and A C Thirlwall

Synopsis

Electrical filters that function simultaneously in two dimensions are not commonly encountered by practising television engineers; but this type of spatial filter plays an important role in the operation of DICE. This section discusses why such filters are needed; the role they play in the conversion process; and their theoretical realisation in digital form. Emphasis is given to the luminance spatial filter and to the manner in which the output is shifted through a quarter line. Spatial filters remove interlace between successive input fields; separate luminance from chrominance; and demodulate chrominance information. Those used in DICE take different proportions of 15 samples from five input points on each of three input lines of the same field.

CLIVE THIRLWALL, BSc, joined the IBA's Engineering Division at Crawley Court on graduating from Manchester University in 1973. He has been primarily concerned with high-speed digital arithmetic subsystems of DICE. He lives in Winchester and has a keen interest in another visual medium – the cinema.

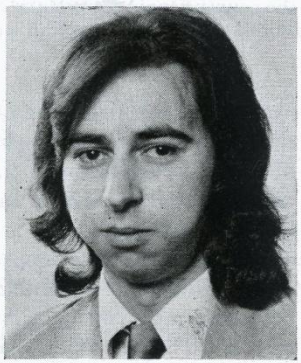

The purpose of the spatial filters used in DICE is to:

(1) Overcome the effects of interlace.
(2) Separate luminance and chrominance information.
(3) In the 625-to-525 direction to re-interlace the signal.
(4) In the 525-to-625 direction to supply 50 lines of signal 'rawstock' per field and to provide half the range of line interpolation.

Why Spatial Filters?

The term 'spatial filter' is not widely used in electronics, the normal electrical filters being, in the television context, one dimensional. A 2.5 MHz low-pass filter will restrict the horizontal resolution and incidentally remove colour, but will not alter the vertical resolution. A vertical aperture corrector will alter vertical resolution, leaving horizontal resolution unchanged.

On the other hand, the spatial filters used in DICE provide a response which is dependent on *both* the horizontal and the vertical frequency of any signal applied to it. By way of example: if a signal of nominal NTSC subcarrier frequency be applied to the luminance spatial filter, but the precise frequency is chosen to be a multiple of line frequency (thus causing the vertical frequency to be zero degrees per line) then, in these particular conditions, the filter provides no attenuation. On the other hand, if the vertical frequency had been π radians or 180° per line (which is the case for true subcarrier) the response would have been zero.

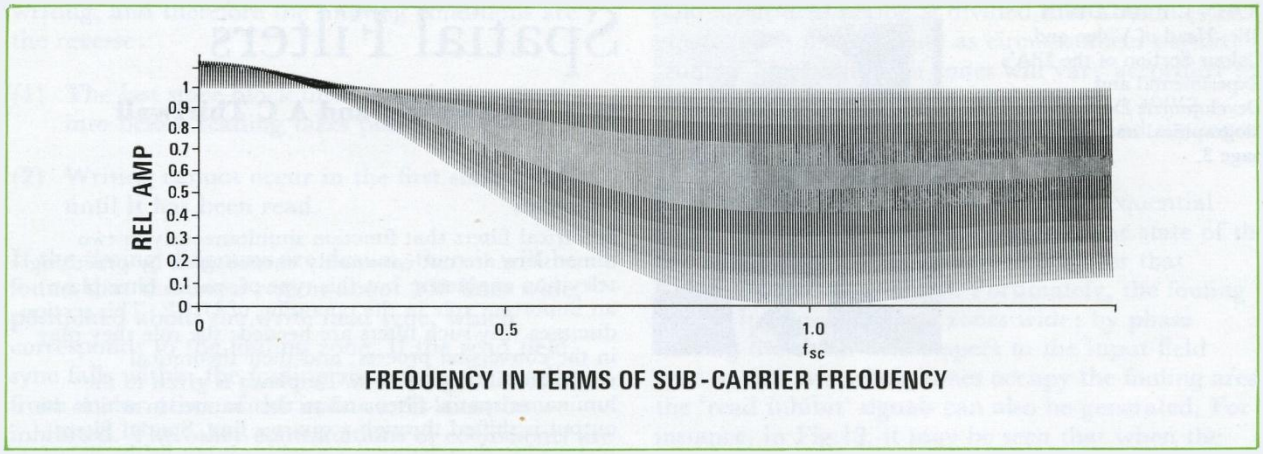

Fig. 1. Frequency response of the luminance spatial filter drawn in the conventional way plotting output against frequency. Even when drawn 1-metre long such representation results in the appearance of spurious shading and would require a high degree of accuracy if accurate information was required.

Figure 1 shows the frequency response of the luminance spatial filter, which rejects chrominance information, drawn in the conventional way. Although the original drawing was approximately 1 metre long, and the increments of the plotter were 0.1 mm (0.004 in), these increments were so coarse that a spurious shading appeared. At about subcarrier frequency, an error of 1 part in 455 in reading the frequency scale would result in the signal appearing to be at a peak in its response rather than at a trough.

Figure 2, on the other hand, shows an isometric view of the response, measured vertically, as a function of horizontal and vertical frequencies. The horizontal frequencies shown are from nominally -5.4 to $+5.4$ MHz and appear along the scale sloping upwards and to the right; whilst the vertical frequencies shown are from $+\pi$ to $-\pi$ radians ($180°$ to $-180°$) per line and these appear along the scale from front to back, sloping upwards and to the left. Beyond these frequency limits the pattern would start to repeat.

From Fig. 2, it can clearly be seen that the horizontal resolution is unimpaired for line repetitive information where the vertical frequency is zero, since the response is constant along the ridge *POQ*. Similarly, the response to vertical luminance frequencies can be seen by considering the ridge *ROS* showing there is a small enhancement of vertical resolution in the spatial filters; this is compensated elsewhere in the converter. However, the diagonal resolution is noticeably affected, since at subcarrier frequency, which corresponds to $\pm\pi$ radians/line and nominally ± 3.6 MHz horizontally, zero responses must occur; only a small response results for frequencies near to this.

Theoretical Realisation

A method of designing spatial filters that has been found effective can be best described as a method of successively meeting different requirements. These filters are digital and all the information that is potentially available is restricted to that occurring at the sample points on a line and on nearby lines.

Figure 3 shows the positions of these samples on odd fields by crosses and on even fields by circles. A requirement of the spatial filters is that they should neutralise the effects of interlace. This is achieved by selecting one vertical position for the output lines and then producing information appropriate for that position from the samples occurring on one input field; and also by producing information appropriate for the same position from the samples occurring on the other input field.

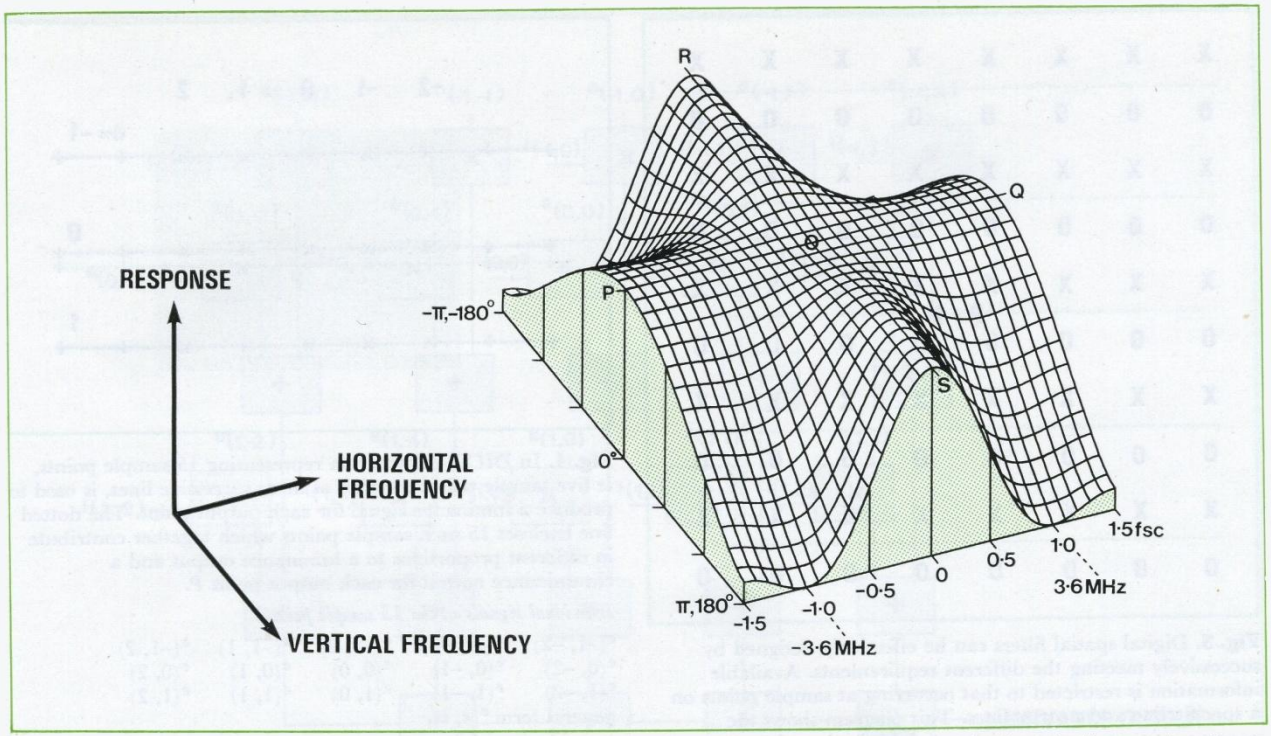

Fig. 2. An isometric view of the response of the luminance spatial filter, and with the response, measured vertically, as a function of horizontal and vertical frequencies. Such a representation is much more informative than the conventional drawing of Fig. 1.

Although any position could be chosen, the one shown*, halfway between the lines of the odd and even fields, is particularly advantageous in that it will result in spatial filters, for the odd and even input fields, which are mirror images of one another. This choice, however, does inevitably cause vertical asymmetry in each spatial filter; this is unfortunate in that it increases the amount of hardware required. Halfway between the lines of the odd and even fields corresponds to a shift of a quarter of the pitch of the lines of a field. On the other hand, there is no restriction in the other direction so each spatial filter can have horizontal symmetry; this symmetry inherently results in zero group delay error in the filter.

In the horizontal direction there are two types of points, exemplified by points P and P' in Fig. 3, which permit horizontal symmetry. Point P is in line vertically with input samples, whereas P' falls halfway between them.

In DICE, information representing 15 sample points, that is five sample points on each of three successive lines, is used to produce a luminance signal and a chrominance signal for each output point; this is shown by the dotted line in Fig. 4 enclosing the 15 input sample points which together contribute in different proportions to a luminance output and a chrominance output for each output point P. If an even number of samples per line had been chosen, the output points would have been positioned as P' in Fig. 3.

The individual signals appearing at the 15 sample points are shown in Fig. 5; any individual signal can be represented by the general term ${}^{e}(n, r)$ by specifying the values of n which may be -1, 0 or 1 and of r which may be -2, -1, 0, 1 or 2.

*In the 625-to-525 direction this state persists during one output field but it alternates with a second symmetrical state during consecutive output fields; this other symmetrical case is with point P immediately below a circle and above a cross. In the 525-to-625 direction the operation is more complex in order to provide roughly a half of the line interpolation; this is achieved by switching from one state to the other every few lines. Whenever this occurs the next vertical output position is arranged to be separated from the previous position by only one-half of the pitch of the lines of a field. This increases the number of lines from 525 to 625 and provides the additional 50 lines per field of signal 'rawstock' required by the line interpolation.

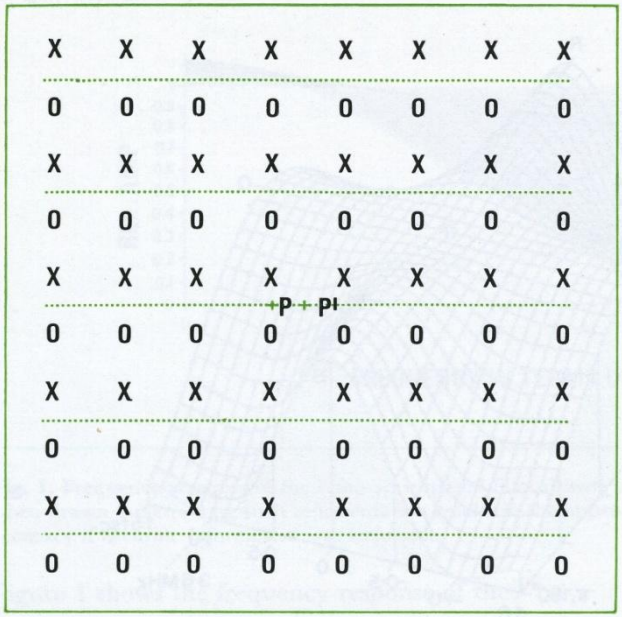

Fig. 3. Digital spatial filters can be effectively designed by successively meeting the different requirements. Available information is restricted to that occurring at sample points on a specific line and nearby lines. This diagram shows the positions of these samples with odd fields indicated by X's and even fields by O's. The position shown, halfway between the lines of odd and even fields, is particularly advantageous.

The luminance spatial filter operates in principle by taking each of the signals $^{e}(n, r)$ multiplying it by its appropriate luminance coefficient $^{l}(n, r)$, and summing the 15 products as shown in Fig. 5.

The chrominance spatial filter operates in principle by performing a similar operation but using the chrominance coefficients $^{c}(n, r)$, (see below Fig. 4) in the place of luminance coefficients $^{l}(n, r)$.

These operations may be expressed mathematically as follows:

Luminance at P

$$L_p = \Sigma\, {}^{e}(n, r) \times {}^{l}(n, r)$$

For $n = -1, 0, 1$
and $r = -2, -1, 0, 1, 2$

Chrominance at P

$$C_p = \Sigma\, {}^{e}(n, r) \times {}^{c}(n, r)$$

For $n = -1, 0, 1$
and $r = -2, -1, 0, 1, 2$

The symbol Σ means the sum of such terms and is the Greek capital sigma.

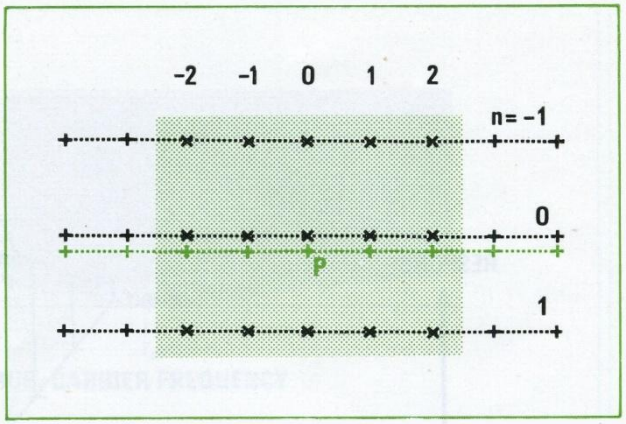

Fig. 4. In DICE, information representing 15 sample points, ie five sample points on each of three successive lines, is used to produce a luminance signal for each output point. The dotted line encloses 15 such sample points which together contribute in different proportions to a luminance output and a chrominance output for each output point *P*.

Individual signals at the 15 sample points:

$^{e}(-1, -2)$	$^{e}(-1, -1)$	$^{e}(-1, 0)$	$^{e}(-1, 1)$	$^{e}(-1, 2)$
$^{e}(0, -2)$	$^{e}(0, -1)$	$^{e}(0, 0)$	$^{e}(0, 1)$	$^{e}(0, 2)$
$^{e}(1, -2)$	$^{e}(1, -1)$	$^{e}(1, 0)$	$^{e}(1, 1)$	$^{e}(1, 2)$

general term $^{e}(n, r)$

Luminance coefficients:

$^{l}(-1, -2)$	$^{l}(-1, -1)$	$^{l}(-1, 0)$	$^{l}(-1, 1)$	$^{l}(-1, 2)$
$^{l}(0, -2)$	$^{l}(0, -1)$	$^{l}(0, 0)$	$^{l}(0, 1)$	$^{l}(0, 2)$
$^{l}(1, -2)$	$^{l}(1, -1)$	$^{l}(1, 0)$	$^{l}(1, 1)$	$^{l}(1, 2)$

general term $^{l}(n, r)$

Chrominance coefficients:

$^{c}(-1, -2)$	$^{c}(-1, -1)$	$^{c}(-1, 0)$	$^{c}(-1, 1)$	$^{c}(-1, 2)$
$^{c}(0, -2)$	$^{c}(0, -1)$	$^{c}(0, 0)$	$^{c}(0, 1)$	$^{c}(0, 2)$
$^{c}(1, -2)$	$^{c}(1, -1)$	$^{c}(1, 0)$	$^{c}(1, 1)$	$^{c}(1, 2)$

general term $^{c}(n, r)$

For line repetitive signals, that is signals which contain only multiples of line frequency (or, expressing this in another way, where the vertical frequency is zero) it is required that the luminance response shall be constant, independent of frequency. This can be realised by making the sum of each vertical column of the luminance coefficients equal to zero, excepting only the central column.

That is:

$$l_{(-1,-2)} + l_{(0,-2)} + l_{(1,-2)} = 0 \qquad (1)$$

$$l_{(-1,-1)} + l_{(0,-1)} + l_{(1,-1)} = 0 \qquad (2)$$

$$l_{(-1,\ 1)} + l_{(0,\ 1)} + l_{(1,\ 1)} = 0 \qquad (3)$$

$$l_{(-1,\ 2)} + l_{(0,\ 2)} + l_{(1,\ 2)} = 0 \qquad (4)$$

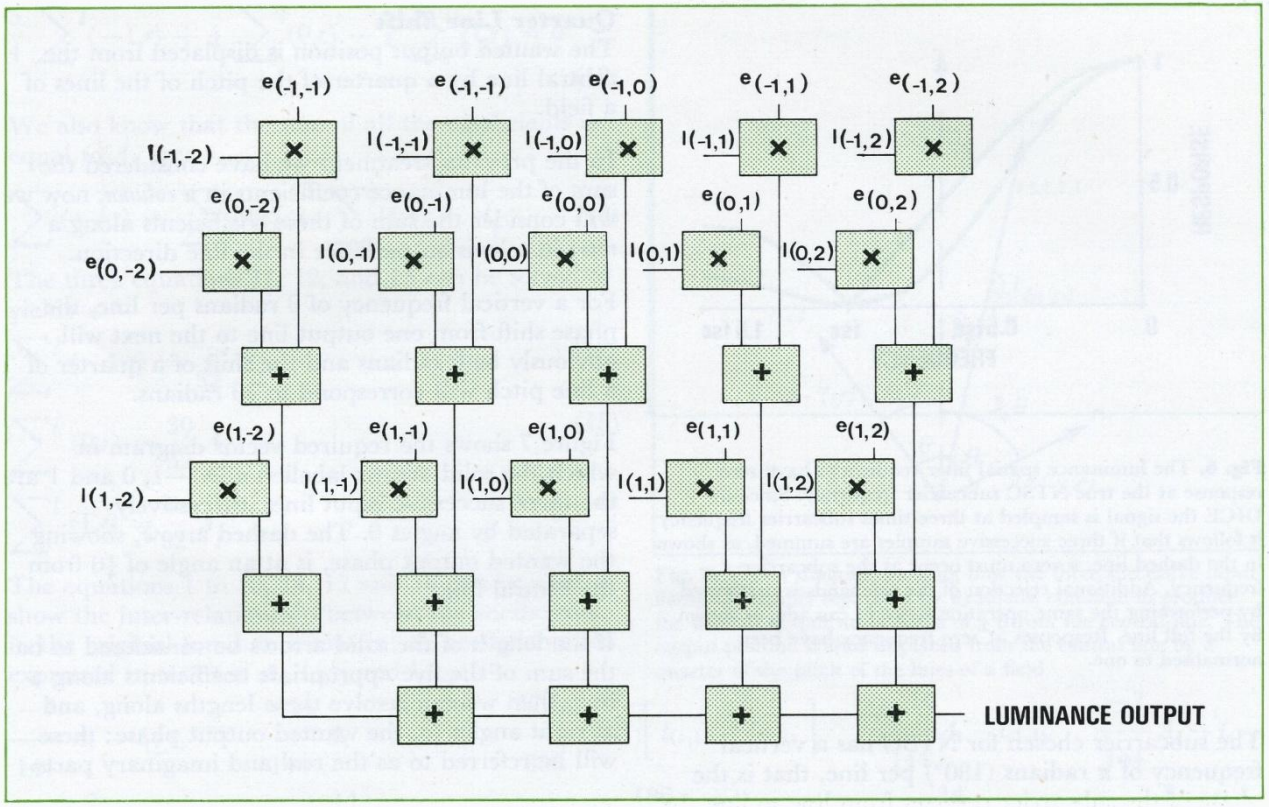

Fig. 5. The luminance spatial filter operates in principle by multiplying each of the sample signals, multiplying it by the appropriate luminance coefficient and then summing the 15 products as indicated in this diagram.

The gain for line repetitive signals, where the vertical frequency is inevitably zero, is given by the sum of all the luminance coefficients; but since all the sums of the vertical columns, save the central one, have already been made zero, then the gain, A_l, is now defined by the central column sum, that is:

$$A_l = l_{(-1,\,0)} + l_{(0,\,0)} + l_{(1,\,0)} \qquad (5)$$

It is necessary that the luminance spatial filter should have zero response at true NTSC subcarrier frequency. In DICE, the signal is sampled at n times subcarrier frequency; it follows that if n successive samples be added together, a zero will inevitably occur at subcarrier frequency. In this case $n = 3$ and the frequency response obtained by summing three successive samples is shown by the upper line in Fig. 9. This gives perfect rejection of subcarrier but insufficient rejection of sidebands. However, this may be overcome by two coincident zeros, that is by performing the same operation twice in cascade.

For the first operation the coefficients of three successive words are:

1, 1, 1

The second operation takes this answer and adds it to the previous two answers thus:

```
          1, 1, 1
             1, 1, 1
                1, 1, 1
Total:    1, 2, 3, 2, 1
```

The two operations in cascade are equivalent to a digital filter in which the coefficients of successive words are 1, 2, 3, 2, 1, and the resulting frequency response is shown by the lower line in Fig. 6; the responses at zero frequency have been normalised to one.

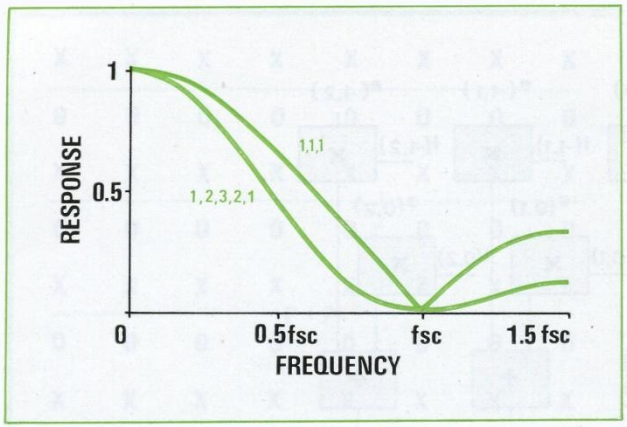

Fig. 6. The luminance spatial filter requires to have zero response at the true NTSC subcarrier frequency. Since in DICE the signal is sampled at three times subcarrier frequency it follows that if three successive samples are summed, as shown in the dashed line, a zero must occur at the subcarrier frequency. Additional rejection of the sidebands is performed by performing the same operation twice in cascade, as shown by the full line. Responses at zero frequency have been normalised to one.

The subcarrier chosen for NTSC has a vertical frequency of π radians (180°) per line, that is the phase of the subcarrier reverses from line to line. In uniform areas of colour, it therefore follows that the chrominance component of an input sample ${}^{e}(\pm 1, r)$ on the top and bottom lines is reversed in phase to the equivalent sample ${}^{e}(o, r)$ on the centre line. By reversing the sign of the coefficients of the top and bottom coefficients before adding the columns of coefficients we can control the horizontal frequency response to signals of vertical frequency equal to π radians (180°) per line; this includes subcarrier frequency. If we arrange that the sums of the columns of the coefficients with signs changed as above, are in the ratio 1, 2, 3, 2, 1, then signals with a vertical frequency of π radians (180°) per line will have a horizontal frequency response as shown by the lower line of Fig. 6.

$$-l_{(-1,-2)} + l_{(0,-2)} - l_{(1,-2)} = v \quad (6)$$

$$-l_{(-1,-1)} + l_{(0,-1)} - l_{(1,-1)} = 2v \quad (7)$$

$$-l_{(-1,0)} + l_{(0,0)} - l_{(1,0)} = 3v \quad (8)$$

$$-l_{(-1,1)} + l_{(0,1)} - l_{(1,1)} = 2v \quad (9)$$

$$-l_{(-1,2)} + l_{(0,2)} - l_{(1,2)} = v \quad (10)$$

where v is a constant

Quarter Line Shift

The wanted output position is displaced from the central line by a quarter of the pitch of the lines of a field.

In the previous treatment we have considered the sum of the luminance coefficients in a *column,* now we will consider the sum of these coefficients along a *row*; in television parlance in the line direction.

For a vertical frequency of θ radians per line, the phase shift from one output line to the next will obviously be θ radians and the shift of a quarter of a line pitch will correspond to $\frac{1}{4}\theta$ radians.

Figure 7 shows the required vector diagram in which the solid arrows labelled $n = -1$, 0 and 1 are the three successive input lines, successively separated by angles θ. The dashed arrow, showing the wanted output phase, is at an angle of $\frac{1}{4}\theta$ from the central line.

If the length of the solid arrows be considered to be the sum of the five appropriate coefficients along a row, then we can resolve these lengths along, and at right angles to, the wanted output phase; these will be referred to as the real and imaginary parts.

Real part

$$\left[\sum{}^{l}(-1,r)\right]\cos\frac{5}{4}\theta + \left[\sum{}^{l}(o,r)\right]\cos\tfrac{1}{4}\theta + \left[\sum{}^{l}(1,r)\right]\cos\tfrac{3}{4}\theta$$

It can be shown that if this be doubly differentiated and equated to zero for θ approaching zero, we obtain a monotonic real part response:

$$\frac{25}{16}\sum{}^{l}(-1,r) + \frac{1}{16}\sum{}^{l}(0,r) + \frac{9}{16}\sum{}^{l}(1,r) = 0 \quad (11)$$

Imaginary part

$$\left[\sum{}^{l}(-1,r)\right]\sin\frac{5}{4}\theta + \left[\sum{}^{l}(0,r)\right]\sin\tfrac{1}{4}\theta - \left[\sum{}^{l}(1,r)\right]\sin\tfrac{3}{4}\theta$$

If this be differentiated and equated to zero for θ approaching zero, we obtain a monotonic imaginary part response which will also be zero at zero vertical frequency:

$$\frac{5}{4}\sum^{l}(-1,r) + \tfrac{1}{4}\sum^{l}(0,r) - \tfrac{3}{4}\sum^{l}(1,r) = 0 \qquad (12)$$

We also know that the sum of all the coefficients is equal to A_l, ie:

$$\sum^{l}(-l,r) + \sum^{l}(0,r) + \sum^{l}(1,r) = A_1 \qquad (13)$$

The three equations 11, 12, and 13 can be solved to yield:

$$\sum^{l}(-1,r) = \frac{-3}{32}A_l \qquad (14)$$

$$\sum^{l}(0,r) = \frac{30}{32}A_l \qquad (15)$$

$$\sum^{l}(1,r) = \frac{5}{32}A_l \qquad (16)$$

The equations 1 to 10, 14, 15 and 16 can be used to show the inter-relationship between the coefficients in the original luminance matrix; all coefficients are expressed in terms of A_l, $l_{(1,1)}$ and $l_{(1,2)}$:

$\frac{-7}{144}A_l - l_{(1,2)}$	$\frac{-14}{144}A_l - l_{(1,1)}$	$\frac{57A_l}{288} + 2\left[l_{(1,1)} + l_{(1,2)}\right]$	$\frac{-14}{144}A_l - l_{(1,1)}$	$\frac{-7}{144}A_l - l_{(1,2)}$
$\frac{7}{144}A_l$	$\frac{14}{144}A_l$	$\frac{186}{288}A_l$	$\frac{14}{144}A_l$	$\frac{7}{144}A_l$
$l_{(1,2)}$	$l_{(1,1)}$	$\frac{45}{288}A_l - 2\left[l_{(1,1)} + l_{(1,2)}\right]$	$l_{(1,1)}$	$l_{(1,2)}$

The denominators of 144 and 288 are most inconvenient in that they are not powers of two. This can be overcome by alterations to equations 14, 15 and 16 as follows:

$$\sum l_{(-1,r)} = \frac{-6}{32}A_1 \qquad (14a)$$

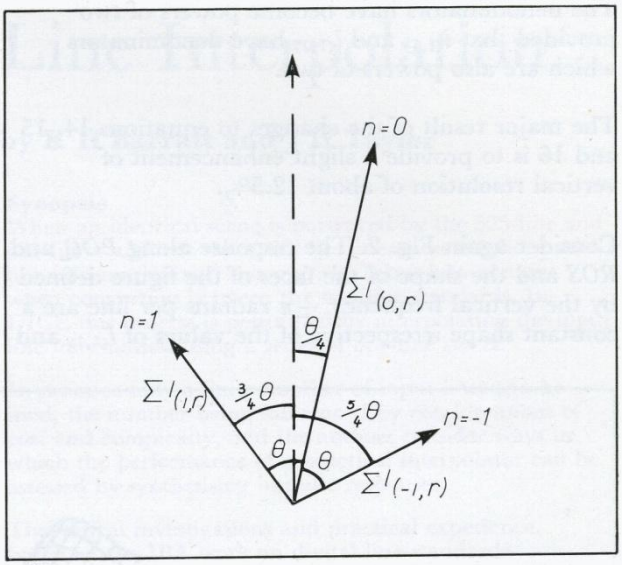

Fig. 7. Vector diagram showing how the three successive input lines, successfully separated by θ radians, combine to produce the wanted output at an angle of $\frac{1}{4}\theta$ from the central-line. The output position is thus displaced from the central line by a quarter of the pitch of the lines of a field.

$$\sum l_{(0,r)} = \frac{34}{32}A_1 \qquad (15a)$$

$$\sum l_{(1,r)} = \frac{4}{32}A_1 \qquad (16a)$$

Using these different equations in place of the previous ones the inter-relationship of the luminance coefficients is now:

$-\frac{A_l}{16} - l_{(1,2)}$	$-\frac{A_l}{8} - l_{(1,1)}$	$\frac{3}{16}A_l + 2\left[l_{(1,1)} + l_{(1,2)}\right]$	$\frac{1}{8}A_l - l_{(1,1)}$	$-\frac{A_l}{16} - l_{(1,2)}$
$\frac{A_l}{16}$	$\frac{A_l}{8}$	$\frac{11}{16}A_l$	$\frac{1}{8}A_l$	$\frac{A_l}{16}$
$l_{(1,2)}$	$l_{(1,1)}$	$\frac{1}{8}A_l - 2\left[l_{(1,1)} + l_{(1,2}\right]$	$l_{(1,1)}$	$l_{(1,2)}$

The denominators have become powers of two provided that $l_{(1,1)}$ and $l_{(1,2)}$ have denominators which are also powers of two.

The major result of the changes to equations 14, 15 and 16 is to provide a slight enhancement of vertical resolution of about 12.5%.

Consider again Fig. 2. The response along *POQ* and *ROS* and the shape of the faces of the figure defined by the vertical frequency $\pm\pi$ radians per line are a constant shape irrespective of the values of $l_{(1,1)}$ and $l_{(1,2)}$. The shape of the faces defined by the horizontal frequency $1.5\,f_{sc}$ is independent of $l_{(1,2)}$ but does depend on $l_{(1,1)}$. Elsewhere the shape depends on both $l_{(1,1)}$ and $l_{(1,2)}$. The gain is controlled by A_l.

Chrominance Spatial Filter

Space does not permit considering the design of the chrominance filter but Fig. 8 shows the shape of the response. Again it is produced by a 5×3 matrix, and has reasonably constant response for frequencies close to subcarrier frequency defined by $\pm f_{sc}$, $\pm\pi$ radians per line.

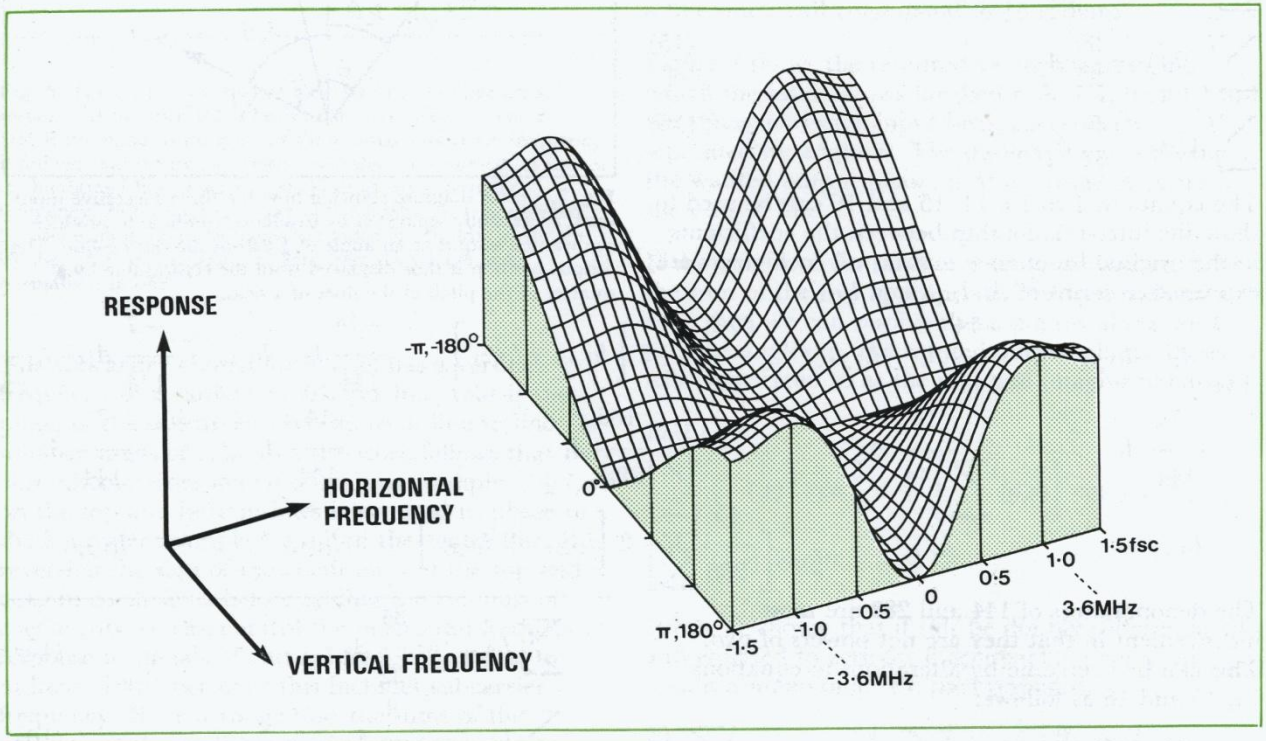

Fig. 8. Shows the shape of the response of the chrominance spatial filter, produced by a five-by-three matrix and designed for reasonably constant response for frequencies close to the subcarrier frequency.

KEN BARRATT, BSc, CEng, MIERE joined the Authority in 1970 and worked initially on the use of digital computers for the monitoring of UHF television transmitters. In 1973 he joined Video and Colour Section as DICE project manager. A graduate of Reading University his previous experience included four years with BBC Designs Department working in the design of uhf television transposers.

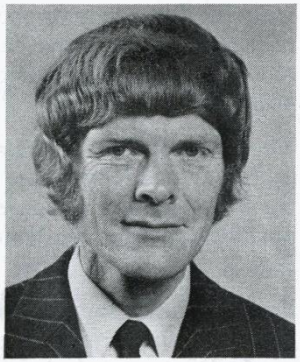

JOHN TAYLOR, CEng, MIERE worked on radar in HM Forces and subsequently in 1960 joined the BBC Designs Department. In 1969 he became a member of the Video and Colour Section of the IBA. Since 1970 has been working on the digital line-rate standards converter and subsequently on DICE.

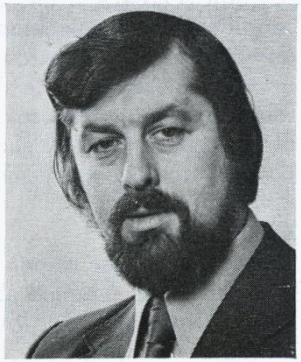

Line Interpolation

by K H Barratt and J H Taylor

Synopsis

When an identical scene is portrayed by the 525-line and 625-line standards there will be differences on the spatial line information and this must be taken into account when converting between the two line standards. In DICE this process is performed by interpolating the input line information using a series of one-line stores.

In practice only a finite number of input lines can be used, the number being influenced by considerations of cost and complexity, and the authors consider ways in which the performance of a practical interpolator can be assessed by synthesising impulse responses.

Theoretical investigations and practical experience, based on the IBA work on digital line-standards converters as well as the first DICE prototype, confirm that good results can be achieved at reasonable complexity, by interpolating five input lines for each output line. This number has been adopted for the final design of DICE and, together with the use of difference interpolation, defines the system organisation. In practice the digital interpolator consists basically of multipliers, adder/subtractors and one-line stores.

The operating requirements are dictated by the data rate and word length. The multiplier used is capable of multiplying an eight-bit data-word by a five-bit coefficient (which is a satisfactory number for this application) to produce 12-bit product at rates well in excess of 18 Mw/s.

The sequential operation of interpolation from line-to-line is organised by line-interpolation control-modules which perform different functions for each direction of operation.

Line interpolation is necessary in standards conversion because the output line structure is required to convey spatial information different from that *directly* available from the input line-structure. This is illustrated in Fig. 1 which shows a sloping line portrayed by a 625-line system (for which the source scanning aperture is infinitely small) and the conversion, without interpolation, to a 525-line structure. The procedure used is equivalent to the discarding of 100 input lines and the uniform redistribution of the remainder.

Recent publications [1, 2] have suggested that interpolation may be regarded as the low-pass filtering of the input line-structure and the resampling of the continuous information, thus obtained, at the required output rate. The conditions required for the perfect recovery of the original scene information are that the low-pass filter should

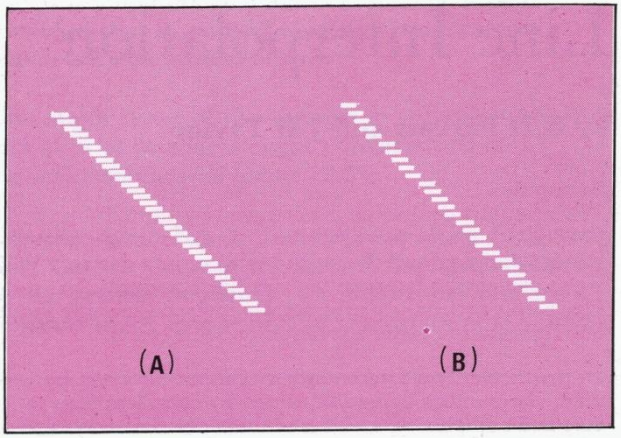

Fig. 1. (*a*) A sloping line as portrayed by a 625-line television raster in which the original scanning aperture is infinitely small. (*b*) The conversion of the sloping line shown in (*a*) without interpolation.

not modify the spatial frequency spectrum of the original scene; the spectrum itself should not exceed one half the input scanning frequency. This leads to a definition of an *ideal* interpolating filter as one having the frequency response characteristic shown in Fig. 2, together with its impulse response.

If the output line rate is less than the input line rate, then either the spectrum of the original scene must be restricted to one half of the output spatial scanning frequency or the low-pass filter must so restrict it. In the latter case however perfect recovery of the original scene information is not possible.

The operation of low-pass filtering may be achieved by the convolution of the input line samples with the filter impulse response*. This leads directly to a possible practical implementation as shown in Fig. 3.

For each input sample, the filter impulse response (scaled to the sample amplitude and centred on the sample) is summed at the required output sampling position (Fig. 3(a)). Fig. 3(b) illustrates a highly simplified practical realisation, in which the one-line delays L ensure that the input sample information is available simultaneously; the multipliers a_1, a_2, etc. scale that information to the filter impulse response shape; and the adder provides the necessary summation.

Limitations

The ideal impulse response shown in Fig. 2(b) is the familiar SINC function and is infinite in extent. Implementation as shown in Fig. 3(b) would require an infinite number of line delays and multipliers. If restricted to a finite number the impulse response

* Sometimes known as the 'aperture function' or 'profile'.

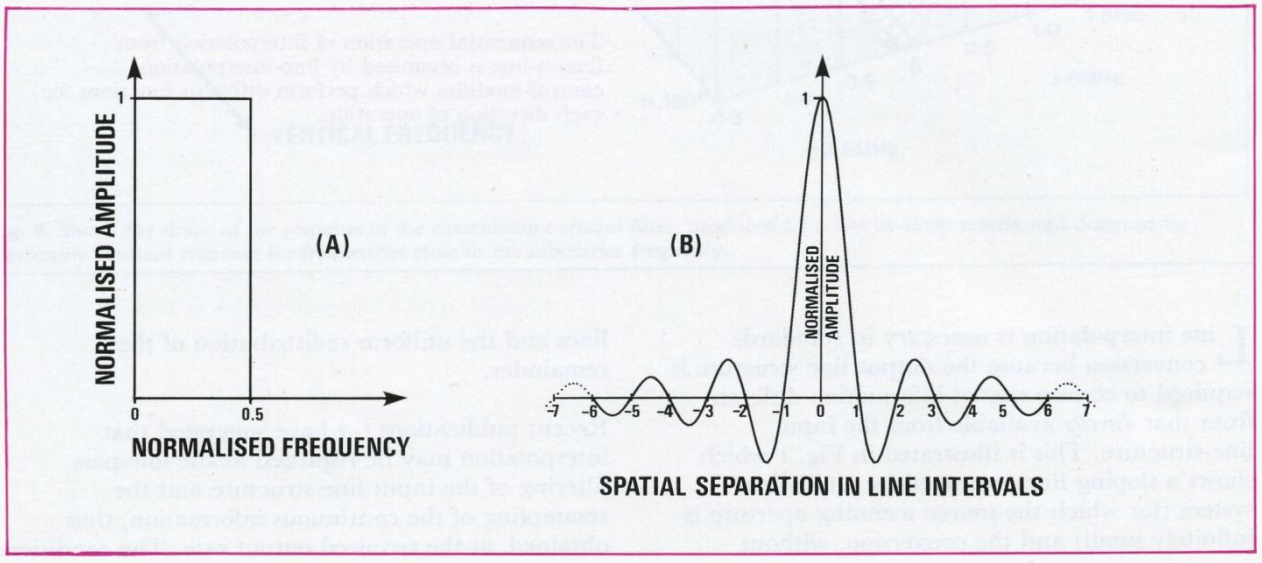

Fig. 2. (*a*) The ideal interpolating filter frequency response which may also be regarded as the spectrum of the impulse response shown in (*b*). (*b*) The impulse response of the ideal interpolating filter which is of the same form as the well-known SINC function.

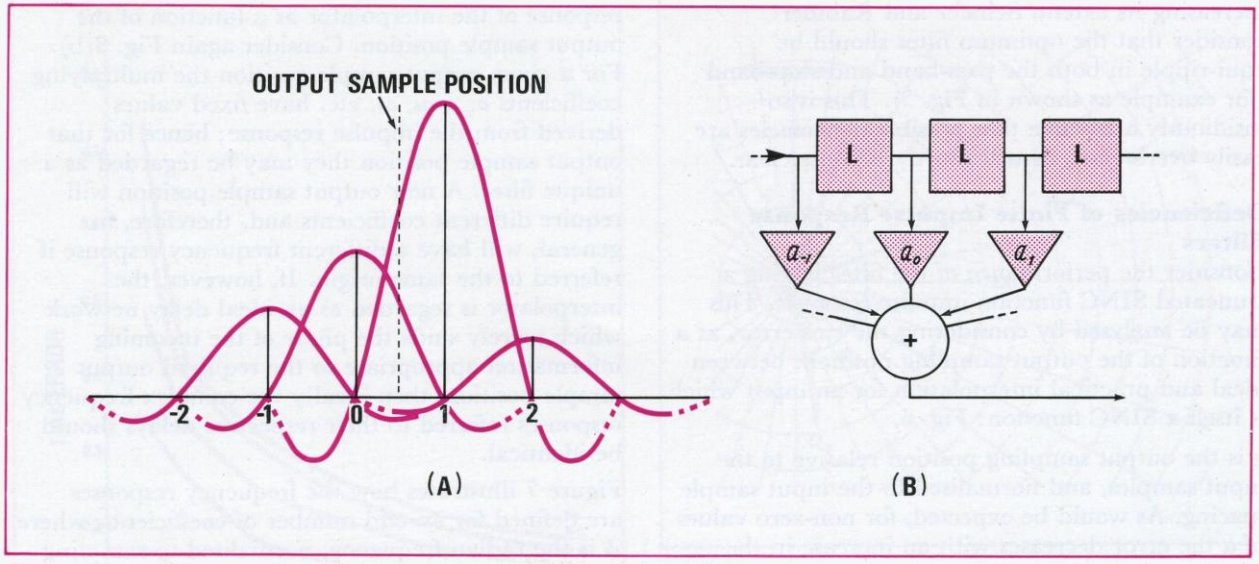

Fig. 3. (*a*) The process of interpolation shown as a convolution procedure which may be directly translated to the practical implementation shown in (*b*).

would be truncated and the low-pass filter would no longer have the ideal frequency response. This is indicated in Fig. 4.

It is not immediately obvious how relatively poor is the performance of the truncated SINC function, nor is it easily seen how to improve it without

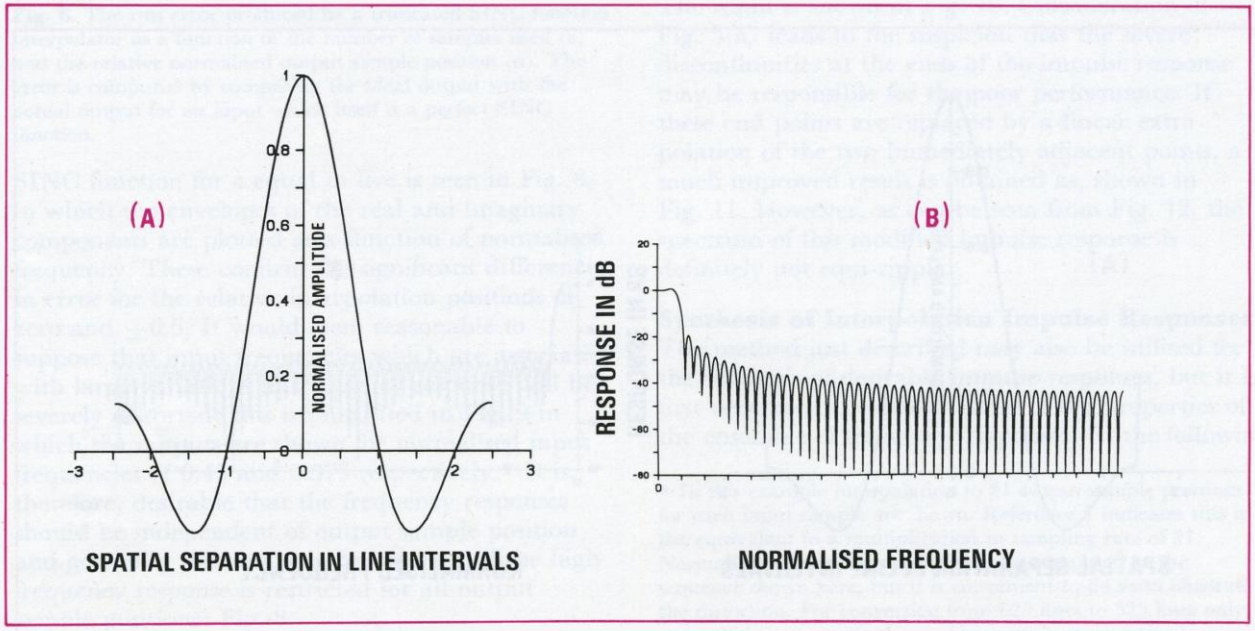

Fig. 4. (*a*) An achievable interpolation filter impulse response consisting of a truncated SINC function which for the example shown here would require five input samples for each output sample. (*b*) The spectrum of the impulse response shown in (*a*) which in comparison with Fig. 2 (*b*) demonstrates its non-ideal nature.

increasing its extent. Schafer and Rabiner[2] consider that the optimum filter should be equi-ripple in both the pass-band and stop-band (for example as shown in Fig. 5). This is so insidiously attractive that possible deficiencies are easily overlooked, as will shortly become clear.

Deficiencies of Finite Impulse Response Filters

Consider the performance of the filter having a truncated SINC function impulse response. This may be analysed by considering the rms error, as a function of the output sampling position, between ideal and practical interpolation for an input which is itself a SINC function: Fig. 6.

a is the output sampling position relative to the input samples, and normalised to the input sample spacing. As would be expected, for non-zero values of a the error decreases with an increase in the number of samples used in interpolation (that is as the impulse response truncation becomes less severe); but, more significantly, for a given value of n, there is a marked variation in error for different values of α.

An alternative method of demonstrating the errors introduced by truncation is to consider the frequency response of the interpolator as a function of the output sample position. Consider again Fig. 3(b). For a given output sample position the multiplying coefficients a_{-1}, a_0, a_1, etc. have fixed values derived from the impulse response; hence for that output sample position they may be regarded as a unique filter. A new output sample position will require different coefficients and, therefore, in general, will have a different frequency response if referred to the same origin. If, however, the interpolator is regarded as an ideal delay network which merely shifts the phase of the incoming information appropriate to the required output sample position, then ideally the complex frequency responses referred to their respective delays should be identical.

Figure 7 illustrates how the frequency responses are defined for an odd number of coefficients, where ϕ is the radian frequency, normalised to sampling frequency.

An error-free interpolator would have a net zero imaginary component for all frequencies, and a constant non-zero resultant along the real axis, for frequencies up to half sampling frequency, at which it would also become zero.

The result of applying this analysis to a truncated

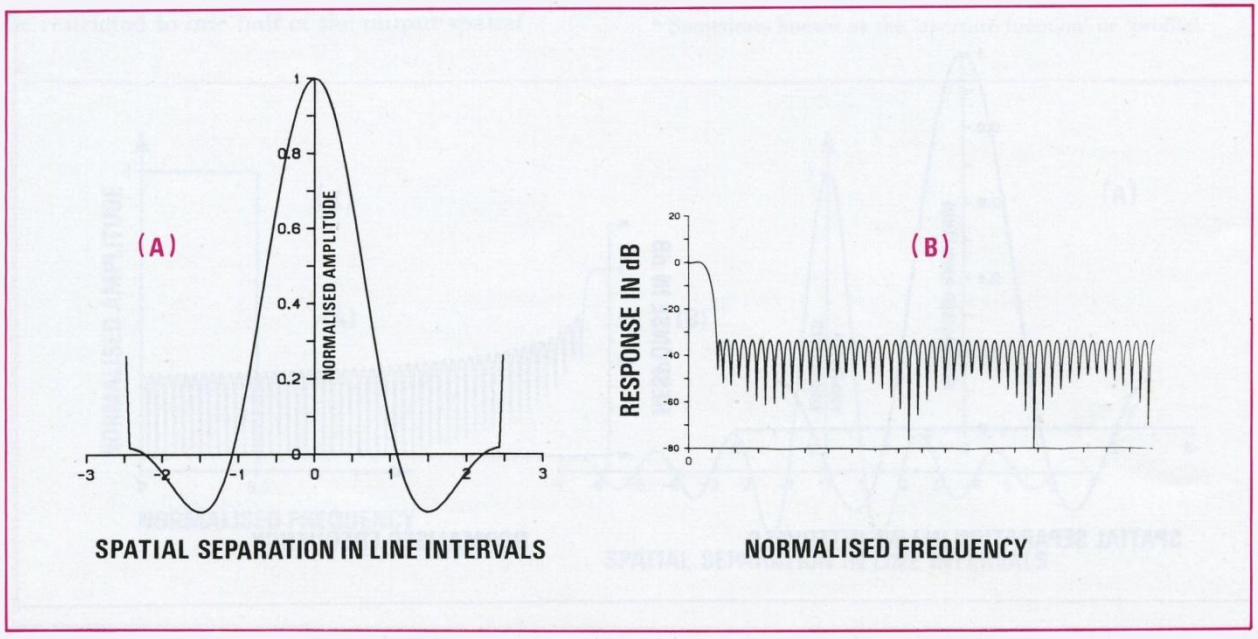

Fig. 5. (*a*) An impulse response obtained by constraining its spectrum shown in (*b*) to be equi-ripple in both pass and stop bands. The severe discontinuities at the ends of the response should be noted.

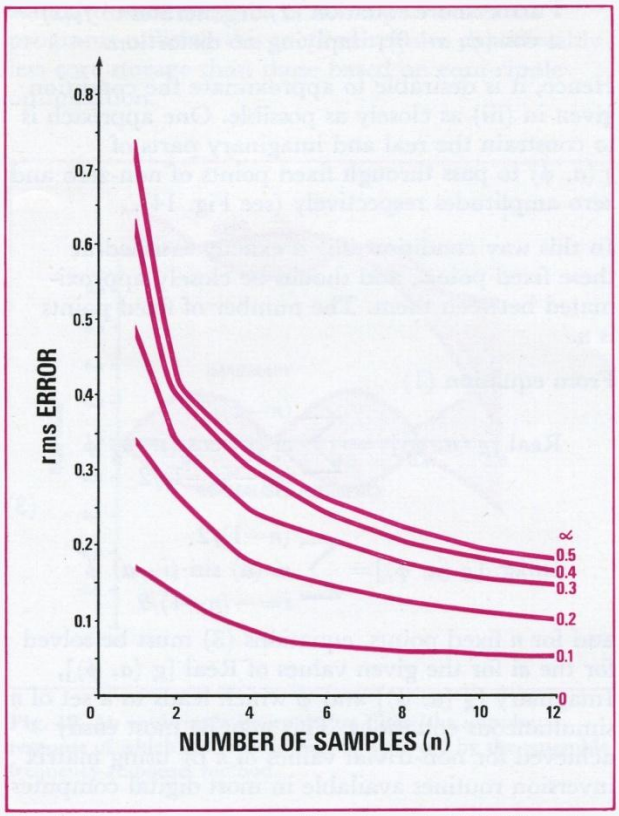

Fig. 6. The rms error produced by a truncated SINC function interpolator as a function of the number of samples used (n) and the relative normalised output sample position (α). The error is computed by comparing the ideal output with the actual output for an input which itself is a perfect SINC function.

SINC function for n equal to five is seen in Fig. 8, in which the envelopes of the real and imaginary components are plotted as a function of normalised frequency. These confirm the significant differences in error for the relative interpolation positions of zero and ± 0.5. It would seem reasonable to suppose that input frequencies which are associated with large variations in frequency response will be severely distorted; this is illustrated in Fig. 9 in which the outputs are shown for normalised input frequencies of 0.45 and 0.375 respectively.* It is, therefore, desirable that the frequency responses should be independent of output sample position and generally this can only be achieved if the high frequency response is restricted for all output sample positions: Fig. 8.

It is of interest to consider the filter shown in Fig. 5, designed using the procedure outlined in

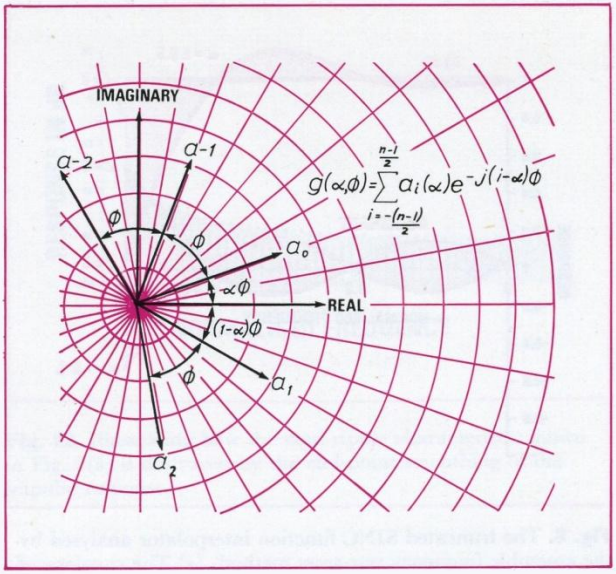

Fig. 7. The definition of the frequency response of an interpolator as a function of output position. ϕ represents the radian frequency normalised to sampling frequency and α is the relative normalised output sampling position. The real axis lies along the desired delayed output position instead of coincident with one of the input vectors which is the usual convention.

reference 1, in the light of the preceding observations. The result is shown in Fig. 10. Consideration of Fig. 5(a) leads to the suspicion that the severe discontinuities at the ends of the impulse response may be responsible for the poor performance. If these end points are replaced by a linear extrapolation of the two immediately adjacent points, a much improved result is obtained as, shown in Fig. 11. However, as can be seen from Fig. 12, the spectrum of this modified impulse response is definitely not equi-ripple.

Synthesis of Interpolation Impulse Responses

The method just described may also be utilised for the synthesis of desirable impulse responses, but it is first necessary to establish some of the properties of the ensemble of frequency responses. In the following

* In this example interpolation to 21 output sample positions for each input sample are shown. Reference 1 indicates this is the equivalent to a multiplication in sampling rate of 21. Normally these positions would not be computed in the sequence shown here, but it is convenient to do so to illustrate the distortion. For conversion from 625 lines to 525 lines only every 25th output sample would be retained (or indeed computed) but eventually all values would appear. In this case, however, 'aliasing' would also have to be considered.

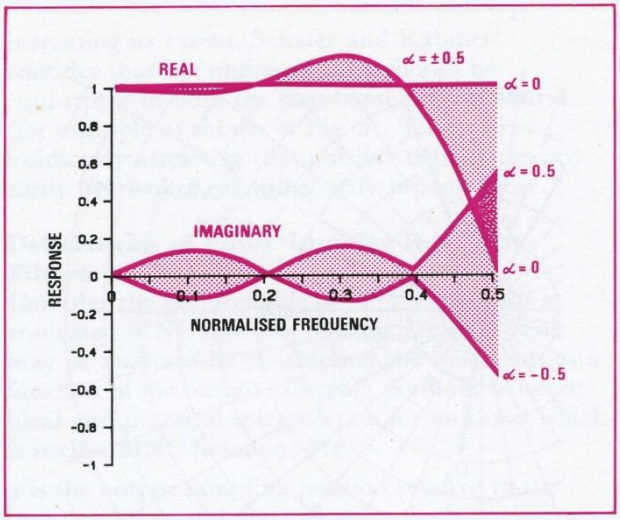

Fig. 8. The truncated SINC function interpolator analysed by the ensemble frequency responses method. (*a*) The envelope of the real part of the responses. (*b*) The corresponding imaginary response envelope.

treatment it will be assumed that an odd number of lines are used for interpolation, though with some modification it is equally applicable to an even number of lines.

The complex frequency response as a function of interpolation position a and normalised input frequency ϕ is given by

$$g\ (a,\ \phi) = \sum_{i=-(n-1)/2}^{(n-1)/2} a_i(\alpha)_e^{\ -j\ (i-a)\ \phi} \quad (1)$$

the relationship between the coefficients $ai(a)$ and the impulse response is shown in Fig. 13.

Now it may be shown that:—

(i) The average of $g\ (a,\ \phi)$ over all a is identical to the impulse response spectrum.

(ii) The output waveform $f_2\ (x)$ for a sampled sine wave input of normalised frequency ϕ_1 and α arbitrary phase θ is given by

$$f_2\ (x) = 2\,|g\ (a,\ \phi_1)|\ \cos\ (\phi_1\ x + \theta + \psi) \quad (2)$$

where $\psi = \tan^{-1}\left[\dfrac{\text{Imag}\ (g\ (a,\ \phi_1))}{\text{Real}\ (g\ (a,\ \phi_1))}\right]$

and a = integer $(x+\frac{1}{2}) - x$

(iii) If the real part of $g\ (a,\ \phi)$ is independent of a and the imaginary part is zero, then the spectrum of the impulse response is zero for normalised frequencies greater than 0.5.

Furthermore equation (2) degenerates to $f_2(x) = A\ \cos\ (\phi_1\ x + \theta)$, implying no distortion.

Hence, it is desirable to approximate the condition given in (iii) as closely as possible. One approach is to constrain the real and imaginary parts of $g\ (a,\ \phi)$ to pass through fixed points of non-zero and zero amplitudes respectively (see Fig. 14).

In this way condition (iii) is exactly satisfied at these fixed points, and should be closely approximated between them. The number of fixed points is n.

From equation (1)

$$\text{Real}\ [g\ (a,\ \phi)] = \sum_{i=-(n-1)/2}^{(n-1)/2} ai\ (a)\ \cos\ (i-a)\ \phi$$

$$\text{Imag.}\ [g\ (a,\ \phi)] = \sum_{i=-(n-1)/2}^{(n-1)/2} ai\ (a)\ \sin\ (i-a)\ \phi \quad (3)$$

and for n fixed points, equations (3) must be solved for the ai for the given values of Real [g $(a,\ \phi)$], Imaginary $[g\ (a,\ \phi)]$ and ϕ which leads to a set of n simultaneous equations. This may be most easily achieved for non-trivial values of n by using matrix inversion routines available in most digital computer

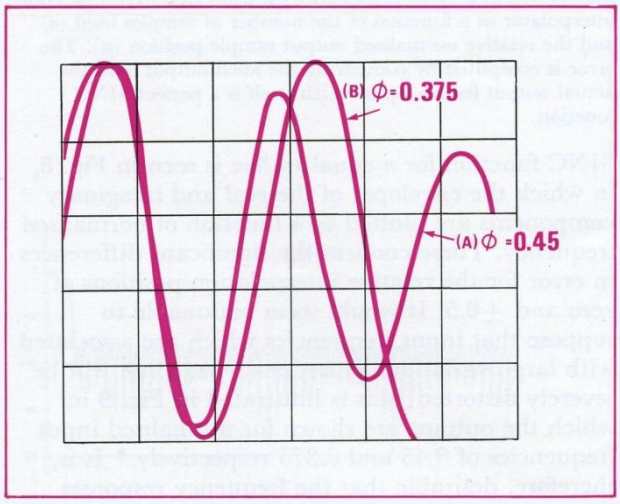

Fig. 9. The result of applying sampled input sinc waves to a truncated SINC function interpolator. Curve (*a*) is the output obtained for a normalised input frequency of 0.45 corresponding to a large variation of frequency responses with a as shown in Fig. 8. Curve (*b*) is the equivalent output for a normalised input frequency of 0.375 for which the variation of frequency responses as shown in Fig. 8 is much reduced.

standard software. It is of interest to note that programs utilising this method require considerably less core storage than those based on equi-ripple optimisation.

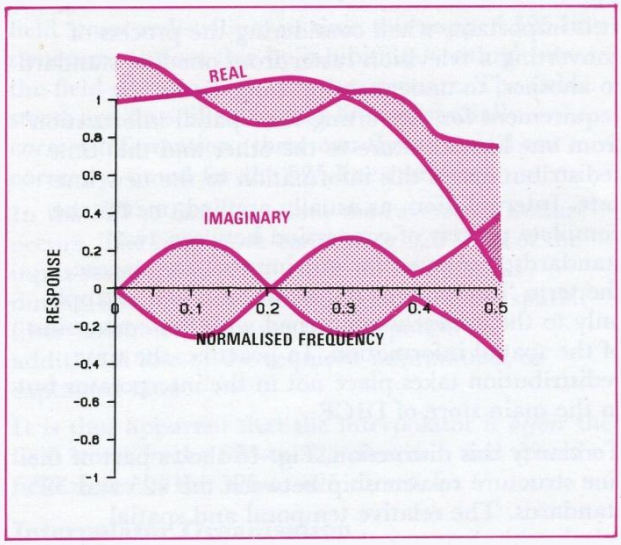

Fig. 10. An equi-ripple interpolation filter (the impulse response of which is shown in Fig. 5) analysed by the ensemble frequency responses method.

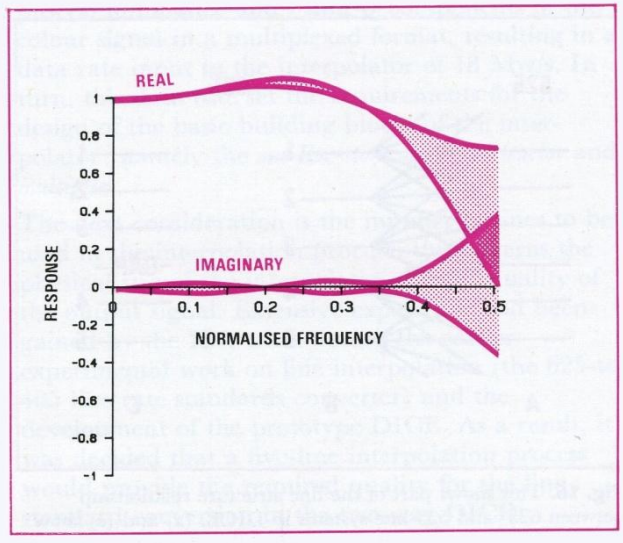

Fig. 11. The greatly improved performance obtained when the impulse response shown in Fig. 5(*a*) is smoothed at the end points.

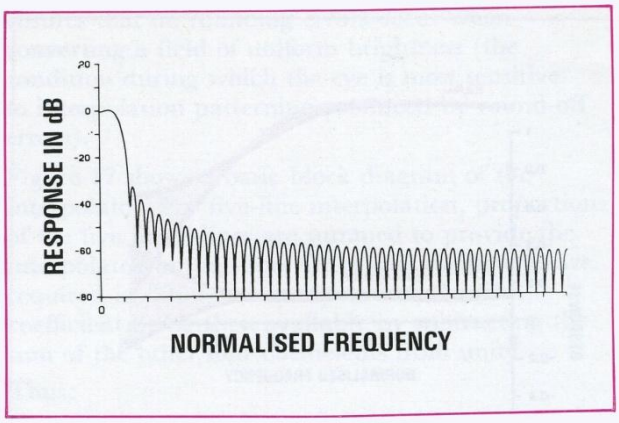

Fig. 12. Illustrating how the equi-ripple characteristic shown in Fig. 5(*b*) is destroyed by the end-point smoothing of the impulse response.

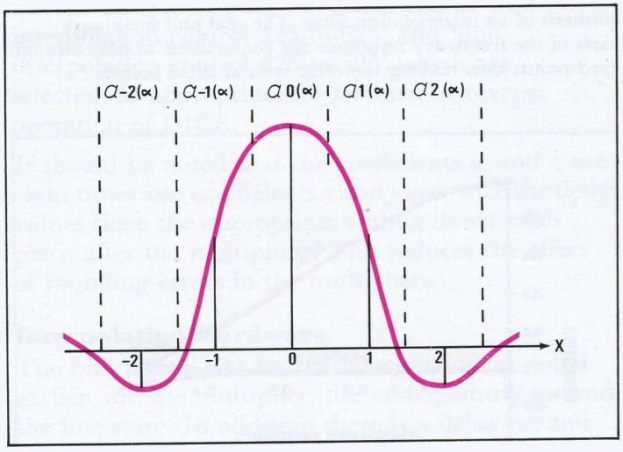

Fig. 13. Demonstrating the relationship between the coefficients $a_i(\alpha)$ and the interpolation impulse response.

Figures 14 and 15 illustrate how high frequency resolution may be traded off against distortion and aliasing.

In summary, it should be noted that there is no generalised practical interpolation filter which is ideal for all purposes, and each case must be considered very much on its merits. For example, the truncated SINC function is not the most suitable impulse response for an interpolator which is used in line-standards conversion; yet, if all input frequencies up to half sampling-frequency were equally likely (and equally important) and if minimum rms error was judged to be the most

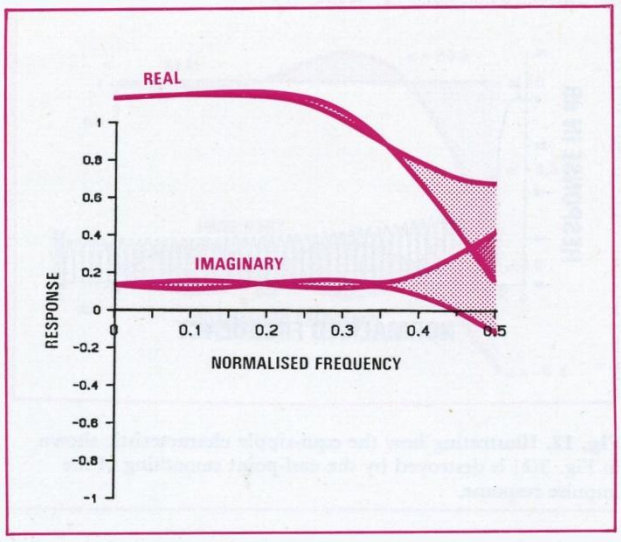

Fig. 14. The ensemble frequency responses method used in the synthesis of an interpolation filter. The real and imaginary parts of the frequency responses are constrained to pass through fixed points thus making the error zero at these points.

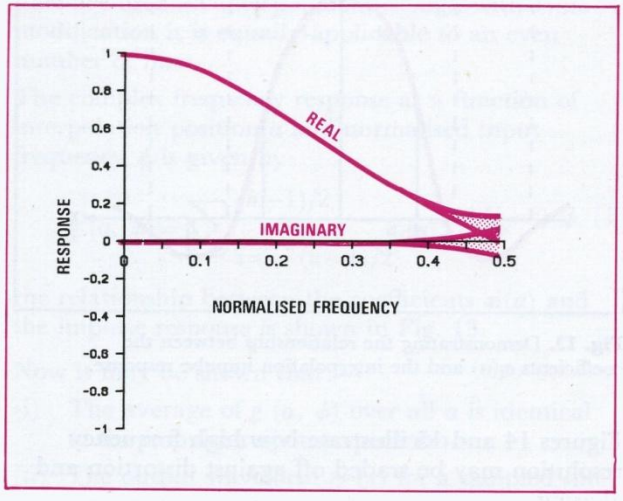

Fig. 15. A further example of the ensemble frequency responses method used in the synthesis of an interpolation filter. In this case the distortion and aliasing rejection performance is enhanced at the expense of high frequency resolution.

important criterion, then it may be demonstrated that such an impulse response is optimum. Furthermore, as has been indicated, optimalisation, defined on the basis of an equi-ripple filtering characteristic, will not necessarily provide the best interpolation strategy. It cannot be too strongly emphasised that the theoretical approach, as described here, *must* be adequately supported by extensive subjective evaluation; nevertheless it can provide a firm basis for investigation.

Practical Considerations

It is important, when considering the process of converting a television raster from one line standard to another, to understand in some detail the requirement for converting the 'spatial information' from one line structure to the other and the 'time redistribution' of this information to the new line rate. Interpolation, as usually applied, means the complete process of conversion between two standards; however, for our immediate purpose, the term 'interpolator hardware' is taken to apply only to the *arithmetic* concerned with the conversion of the spatial information. In practice, the time redistribution takes place not in the interpolator but in the main store of DICE.

To clarify this distinction, Fig. 16 shows part of the line structure relationship between the 625 and 525 standards. The relative temporal and spatial relationship between the structures is shown in Fig. 16(a) and (b). The spatial information for the new line structure (b) is obtained by using information from *five* input lines. However, the information is produced by the interpolation arithmetic at the original line-rate of 625 lines, with

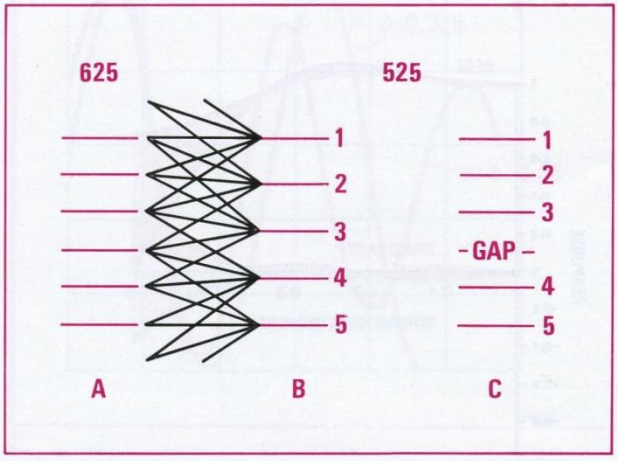

Fig. 16. This shows part of the line structure relationship between 625- and 525-line systems in DICE. (*a*) and (*b*) show the *spatial* relationship between the 625-line raster (one field) and the 525-line raster, indicating by means of black lines the five lines used to interpolate an output line. The relative *time* at which these lines are generated is shown in (*c*).

the correct spatial information on 525 of these lines. Thus, there are 100 lines in every frame which must be discarded by the time redistribution process. These are the so-called *gap* lines shown in Fig. 16(c). There is zero information on these gap lines (designed into the interpolation process) and the field stores redistribute in time this gapped 525-line structure. It does this by inhibiting 'writing' into the field store each time a gap appears. The field stores are thus filled by 525 lines of spatially correct information; these can then be read out correctly, timed for the 525-line standard.

In the 525-to-625 direction, the reverse procedure occurs. The field store containing 525 lines of the input standard *ceases* to read out at intervals; during this time, information stored in the spatial filter 'arithmetic' is used to interpolate an additional line of the required information, as explained later.

It is thus apparent that the interpolator is *before* the field stores in the 625-to-525 direction and *after* the field stores in the 525-to-625 direction.

Interpolator Organisation

An ideal interpolator would require hardware of infinite dimensions. As a more practical example, the design of the line interpolator for DICE was based on many factors: the two main considerations were speed and number of lines of interpolation.

The speed requirement was dictated by the need to process luminance and I and Q components of the colour signal in a multiplexed format, resulting in a data rate input to the interpolator of 18 Mw/s. In turn, this data rate set the requirements for the design of the basic building blocks of the interpolator: namely the *one-line-store, adder/subtractor* and *multiplier.*

The next consideration is the number of lines to be used in the interpolation process: this governs the physical size of the interpolator and the quality of the output signal. Extensive experience had been gained by the IBA team during the earlier experimental work on line interpolation (the 625-to-405 line rate standards converter) and the development of the prototype DICE. As a result, it was decided that a five-line interpolation process would provide the required quality for the line standard conversion in the two-way DICE.

The organisation of the interpolator is in the form known as *difference interpolation.* This allows the use of $(n-1)$ multipliers in an n-line interpolator, and also ensures that no rounding errors occur when converting a field of uniform brightness (the condition during which the eye is most sensitive to interpolation patterning produced by round-off errors).

Figure 17 shows a basic block diagram of the interpolator. For five-line interpolation, proportions of the five input-lines are summed to provide the interpolated output-line. Only four coefficients are required in difference interpolation, the fifth coefficient being then available by subtracting the sum of the other four coefficients from unity.

Thus:

$$\begin{aligned}\text{Output} &= R+x(Q-R)+y(S-R)+w(P-R)+z(T-R)\\ &= R+xQ-xR+yS-yR+wP-wR+zT-zR\\ &= R-R(w+x+y+z)+wP+xQ+yS+zT\\ &= R[1-(w+x+y+z)]+wP+xQ+yS+zT\end{aligned}$$

The coefficients w, x, y, z, required by the interpolator, are under the command of 'line interpolation control 525-to-625' and 'line interpolation control 625-to-525' which are selected, as appropriate for forward or reverse operation of DICE.

It should be noted that the coefficients w and z are eight times and coefficients x and y are twice the design values since the appropriate scaling down takes place after the multipliers. This reduces the effect of rounding errors in the multipliers.

Interpolation Hardware

The building blocks for the interpolator, as noted earlier, are the multiplier, the adder/subtractor and the line store. In addition there is a delay module used in several places to ensure that data passing through the various paths of the interpolator are kept in step. The hardware operates at a data rate of 18 Mw/s.

A major design consideration were the round-off errors. Although the data word is initially of fixed length, the multipliers and adders can cause the word length to grow unmanageably. It was determined that, since the data word is reduced after interpolation to 8 bits by truncation, the word length within the interpolator could be safely limited to 12 bits without unduly increasing quantisation noise by round-off errors.
Another problem stemming from round-off errors was to decide the required accuracy of the coefficient number since this would have a significant bearing on the design of the multiplier.

Again the earlier experience of the IBA team provided guidance but in addition computer simulation of the interpolator was used to determine likely performance with given quantisations of the coefficient number. Finally, it was decided that the multiplier should be based on an 8-bit data word by 5-bit coefficient giving a product truncated to 12 bits.

The Multiplier

Although many ways of designing digital multipliers exist in theory, most practical designs are governed by application requirements, and—in the case of DICE—by the limitations of currently-available integrated circuits. The basic requirement, as indicated above, was an 8-bit by 5-bit multiplier giving a 12-bit product; further since the multiplier is used in a *difference path* (see Fig. 17) it could be presented with *negative* binary numbers (2's complement) from the subtractors. It was thus desirable (though not essential since the problem could have been solved another way) for the data input to the multiplier to accept an 8-bit signed (2's complement) number. Further, since the coefficients could be negative, the 5-bit input should also be able to accept negative numbers; this was found to be difficult to achieve in practice and the difficulty was overcome by turning the adder following the multiplier in the interpolation arithmetic into a subtractor whenever negative coefficients are required. Finally, as the multiplier operates on continuous 18 Mw/s data the design had to be capable of accepting data and producing answers at this rate, though there is no limitation (within reason) on the processing time.

Figure 18 shows a simplified block diagram of the multiplier used in DICE. This indicates that the 5-bit multiplier-word is extended to include 10 0000 so that 32 times the input can be provided. Although the clocking arrangements are not shown in Fig. 18, it should be noted that the process from input to output requires two clock periods. To achieve full accuracy of the product, 14-bits-plus-sign are required; however this is truncated to 12-bits-plus-sign.

The algorithm is based upon a special shift-and-add type of multiplier; since in the binary system each

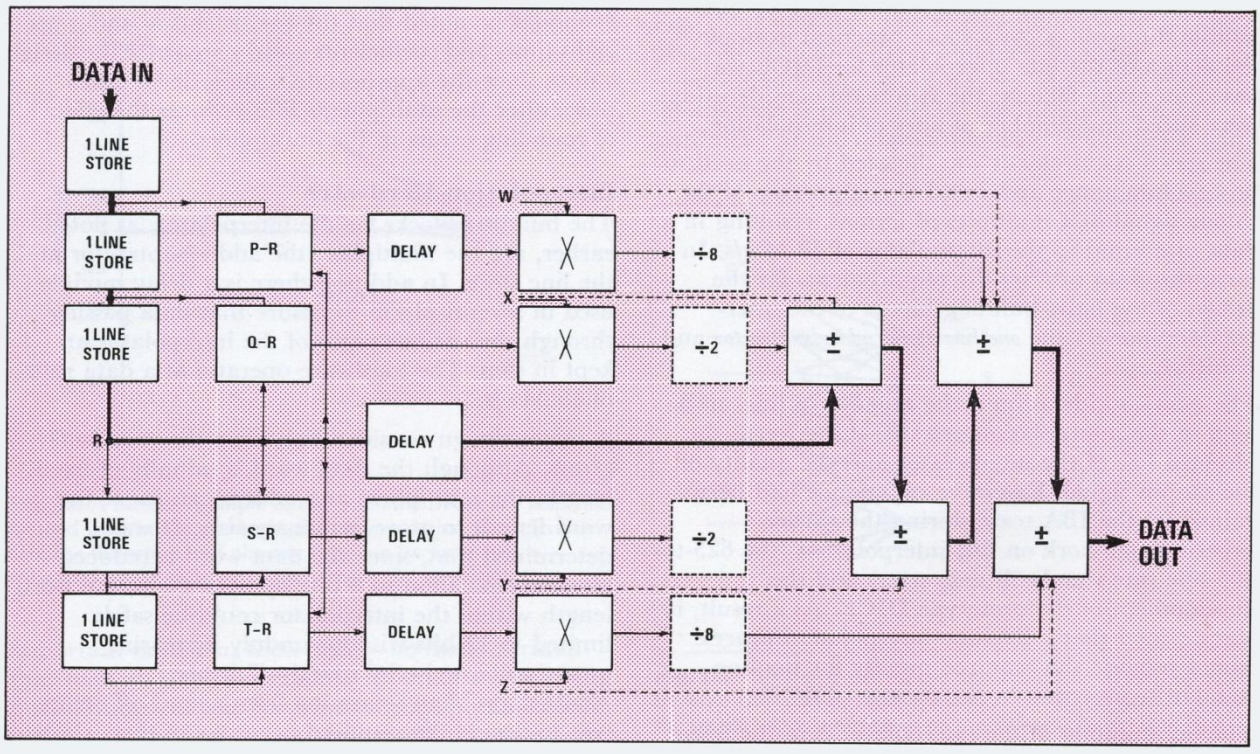

Fig. 17. Simplified block diagram of five-line difference interpolator. Four multipliers are used to proportion five lines using the following equation: output $= [1-(w+x+y+z)]\ R + wP + xQ + yS + zT$.

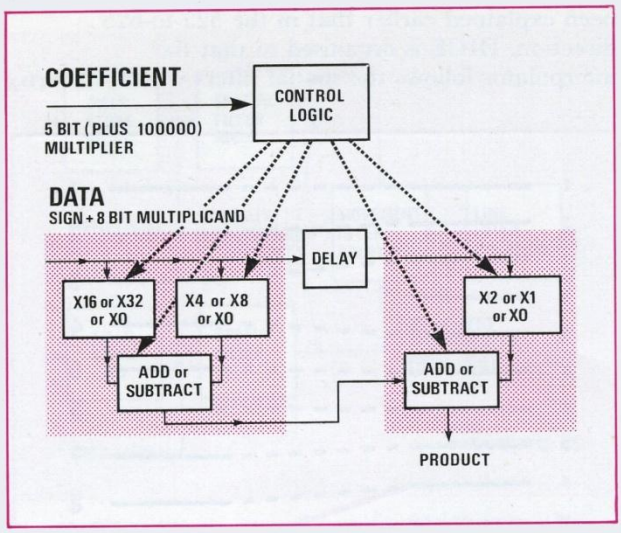

Fig. 18. Simplified block diagram of the multiplier used for movement interpolation and line interpolation. Two stages of addition/subtraction are used to provide a sign plus 8-bit by 5-bit multiplier. In addition, provision is made for the multiplier to be 10 0000 to give unity times the input multiplicand. The multiplier is designed to handle a data rate of 18 Mw/s and a coefficient rate of 16 kw/s.

progressive shift to left is equivalent to multiplying by 2, 4, 8 and so on, any multiplier function may be built up by adding or subtracting an appropriately-shifted input data number. Three simple examples of this operation are:

$$\times\ 3 = \times(2^2 - 2^0) = \times(\ 4 - 1)$$
$$\times 15 = \times(2^4 - 2^0) = \times(16 - 1)$$
$$\times 18 = \times(2^4 + 2^1) = \times(16 + 2)$$

In practice, any multiplying integer, m, in the range 0 to 42 may be realised, as follows:

$m = (32 \text{ or } 16 \text{ or } 0) \pm (8 \text{ or } 4 \text{ or } 0) \pm 2\ (2 \text{ or } 1 \text{ or } 0)$

This expression is equivalent to:

$m = (2^5 \text{ or } 2^4 \text{ or } 0) \pm (2^3 \text{ or } 2^2 \text{ or } 0) \pm (2^1 \text{ or } 2^0 \text{ or } 0)$

To realise the multiplication (Fig. 18), it is necessary only to have two adder/subtractors with arrangements to feed into one port of one of them, data shifted left by either five or four places, or replaced by zeros. Into the other port of this adder/subtractor the data is shifted to the left by either two or three places, or replaced by zeros. The output of this adder feeds one input port of the other adder/subtractor; to the second input port of this second adder/subtractor is fed data either shifted by one place, or unshifted, or replaced by zeros.

To consider Fig. 18 more fully it is necessary to describe the right hand block, marked $\times 2$, $\times 1$ or $\times 0$. A $\times 0$ condition would replace the input data number by 'all zeros'; a $\times 1$ condition would transfer the data directly into the add or subtract block; a $\times 2$ condition shifts the data up one significant bit before transferring the result into the add or subtract block.

The input multiplier number sets up the multiplier via the control logic: if, for example, it were desired to multiply by 29, the binary number representing 29 would (via the control logic) set the first stage to: $\times 32$ subtract $\times 4$. The output from the first stage would then be $\times 28$. The second stage would be set: add $\times 1$. This combination would then give the required output product: $\times 29$.

The 'delay' keeps the data presented to the second stage in step with the output from the first stage.

The One-line Delay

The one-line-delay element in the interpolator (known as an 18 MHz *YIQ* store) can store one complete active television line of data. As already described, the data processed by the interpolator is multiplexed Y, I and Q components in a word format of $Y\ I\ Y\ Q\ Y$. This requires storage of 956 8-bit words.

Because of the high speed of operation, the *dynamic shift-registers* suitable for use as the storage element were incapable of storing input data directly at the 18 MHz rate. The solution was to demultiplex each bit of the 8-bit word into four data streams at a rate of 4.5 MHz each, store these data streams, and subsequently to multiplex the output back to 18 MHz.

The design of the 18 MHz *YIQ* store-board is such that it is able to store only 4 bits of an 8-bit time-multiplexed Y, I and Q signal so that in practice two identical store boards are used to handle the full 8 bits.

In the interpolator, a total of ten *YIQ* stores are required to delay five lines of information. A counted clock-generator is used to ensure that exactly the right number of clock pulses are provided to operate the dynamic shift registers of the stores.

Line Interpolation Control

It was necessary, before calculations affecting the performance of the interpolation process could be considered in relation to the generation of the

interpolation coefficients, to study the input/output line relationships. The interpolation arithmetic, although remaining a processing 'black box', is positioned differently in the organisation of the converter when operated in the 525-to-625 and 625-to-525 modes. DICE has two completely independent line interpolation control systems and the explanation can be made clearer by treating separately the two modes of control.

625-to-525 Interpolation

In this mode, interpolation takes place before the main store. The ratio 625:525 reduces to 25:21 so that there is a complete cycle of input/output line relationships for each 25 lines of input. The spatial relationship for the two line structures is shown in Fig. 19. The dotted lines connecting the two structures show the *nearest* line on the 625-line structure for a given 525 line, indicating the degree of interpolation required to generate the new line. This nearest line is also the *centre* line of the five lines used for the interpolation process: an example of the way in which line 9 of the output structure is produced is also indicated.

Figure 19 indicates that there are four lines in the input structure that do not occur as the nearest line for the output structure. It is at these periods that the output from the line interpolation is not written into the field stores and, as already noted, this is referred to as a *gap*. It will become apparent that although these gaps occur this does not imply that *any* input lines are ignored.

The 25:21 sequence is repeated continuously and completes 25 cycles in any one frame of television picture. The phase of the interpolation sequence is checked at the start of each output field and is reset should it be found out of step. This procedure minimises vertical perturbations of the output picture where a 'gen-lock' is taking place on the output reference sync pulses.

The sequence is under the control of a 25-step-counter which is stepped on by the line sync pulses of the input standard. The output from the counter addresses read-only-memory devices, which store the numerical information of the coefficients for the line interpolation.

525-to-625 Interpolation

This mode is more complex than in the 625-to-525 direction since the interpolation process includes functions carried out by the spatial filters. It has been explained earlier that in the 525-to-625 direction, DICE is organised so that the interpolator follows the spatial filters which, in turn,

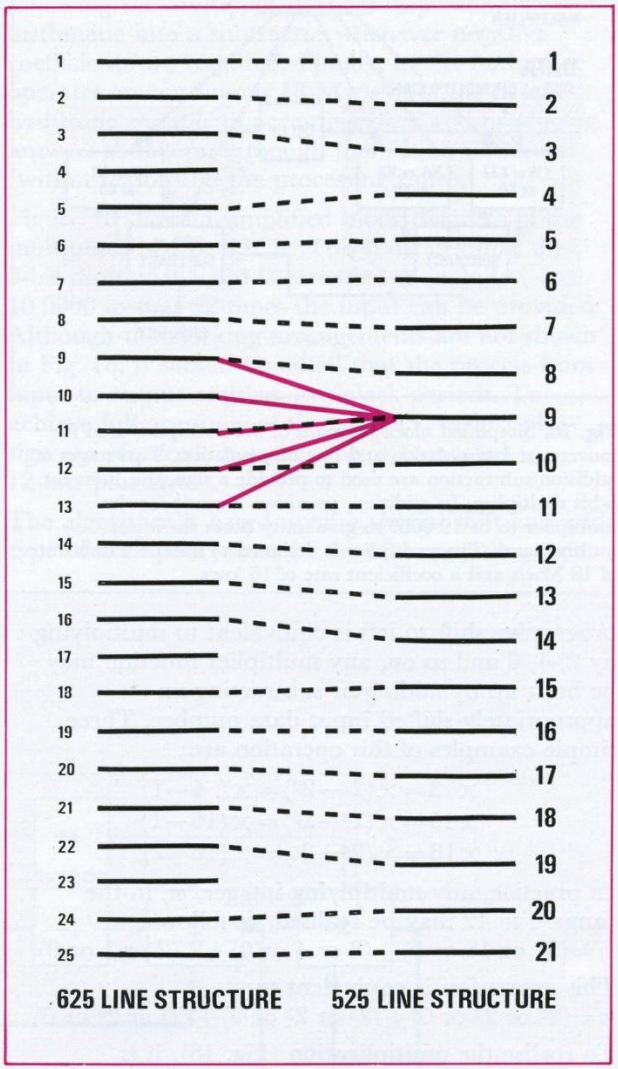

Fig. 19. This shows the spatial relationship between the 625 field structure and the 525 field structure over 25:21 lines (the repeating sequence). Using this diagram, the nearest line on the input structure to the required output line may be found, which will define the centre line of the five lines used in the interpolation process. An example is shown in colour where line 9 on the output structure has line 11 on the input structure as the nearest. This is used, together with lines 9, 10, 12 and 13 as the five lines for the interpolation process. In the sequence shown, lines 4, 10, 17 and 23 are not used as centre lines for the interpolation process. At these points the main store (which follows the interpolator in the 625-to-525 direction) waits for a period of one line.

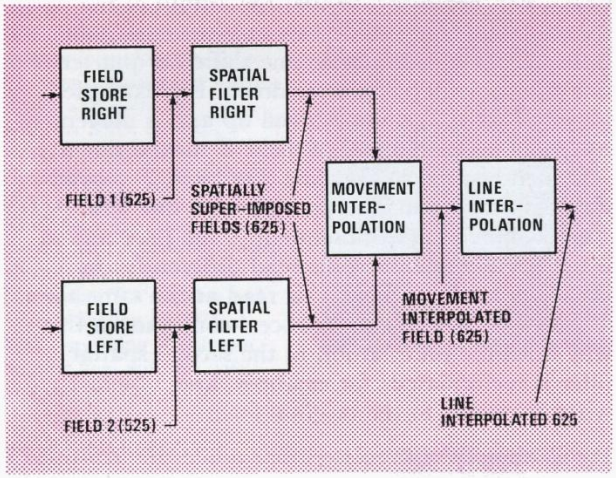

Fig. 20. 525-to-625 direction of operation, showing the relative position of processing associated with line interpolation.

follow the field store. Fig. 20 shows a simplified block diagram of this section of the converter.

The following is a simplified description of operation in this mode and further information is contained in the papers on field stores, spatial filters and movement interpolation. Here we are concerned primarily with the interaction of the circuit blocks shown in Fig. 20.

The field stores normally contain, at any given time, an odd and an even field. The spatial filters interpolate in opposite directions (the quarter-line-shift), thereby producing a de-interlaced raster, or in other words two fields superimposed spatially. This enables the two fields to be mixed in the proportions required for movement interpolation. As the spatial filters can interpolate one-quarter of a line up or down, there is capability of generating spatially coincident lines in any of 525 positions from each of the two original interlaced rasters. Re-interlacing occurs by choosing every other line position for one output field and the intermediate positions for the interlaced output field. In addition, *extra* lines are generated within the spatial filters by recirculation of the data and by changing the direction of the quarter-line shift. In this way, the extra lines of information are generated to create 625 lines from the 525 lines stored in the main store. All the generated lines, however, are spatially related to the input raster by one-quarter of a line shift. It is now necessary to interpolate spatially this information to the correct positions dictated by the 625-line output raster requirements: this occurs in the line interpolator.

Figure 21 shows the spatial relationships following through from the 525-line interlaced raster to a *single* 625-line field. Only one cycle of the 25:21

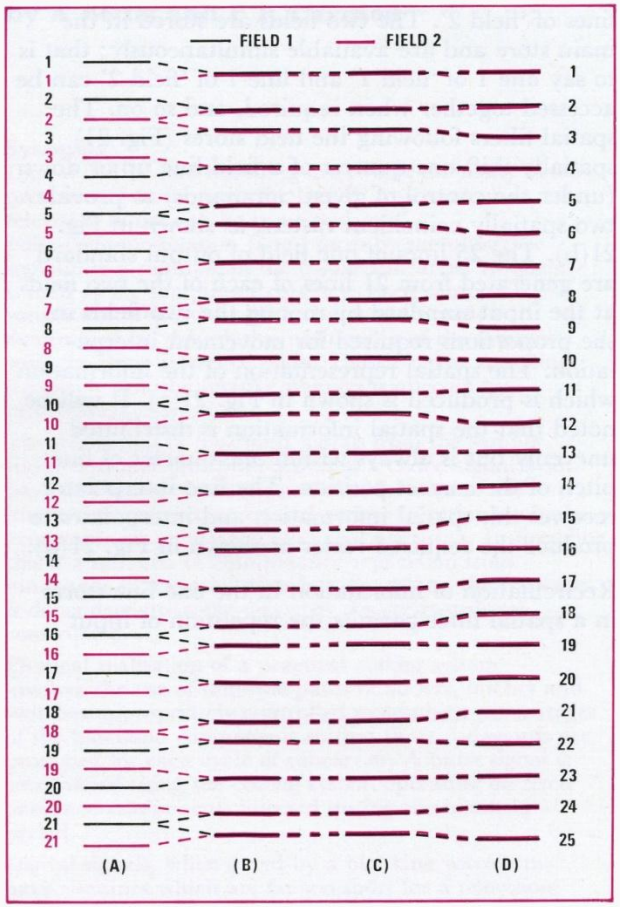

Fig. 21. 525-to-625 direction of operation, showing the spatial relationship of the line structures through the processing shown in Fig. 20. A complete cycle of 21:25 events is shown. (*a*) shows 2 × 21 lines of interlaced fields 1 and 2 stored in the main store. These are accessed in parallel and are interpolated $\frac{1}{4}$-line up or down in the spatial filters to produce 2 × 25 lines of spatially superimposed information. (*b*) shows the two fields now mixed in the proportions required for movement interpolation to produce (*c*). The line interpolator next spatially redistributes (*c*) to produce the correct 625 raster (*d*). (*d*) is shown to be 25 lines of one particular field of the output standard. The interlaced field of the output standard is generated by changing the direction of the $\frac{1}{4}$-line shifts—this is not shown in the diagram.
Where an input line from the main store is used twice—for example, line 2 of 'field 2', the data in the spatial filter line stores is recirculated and therefore made available for re-use.

sequence is shown, but since this cycle is continuously repeated, the diagram indicates all relationships which occur in practice.

On the left of the diagram (a) is shown a spatial relationship of part of a 525-line interlaced raster, with 21 lines of 'field 1' distinguished from the 21 lines of 'field 2'. The two fields are stored in the main store and are available simultaneously; that is to say line 1 of 'field 1' and line 1 of 'field 2' can be accessed together when required, and so on. The spatial filters following the field stores (Fig. 21) spatially shift one-quarter of a field line up or down (under the control of given commands) to provide two spatially coincident rasters, as shown in Fig. 21(b). The 25 lines of one field of output standard are generated from 21 lines of each of the two fields at the input standard by mixing the two fields in the proportions required for movement interpolation. The spatial representation of the information which is produced is shown in Fig. 21(c). It will be noted that the spatial information is distributed unevenly but is always within one-quarter of line pitch of the correct position. The line interpolator receives this spatial information and interpolates to produce the required raster as shown in Fig. 21(d).

Recirculation of information in the one-line stores in a spatial filter permits the repetition of input lines when necessary, so that the output of the spatial filter can provide two successive outputs from the same input data; one shifted a quarter of a line up, the other a quarter down. For example, line 2 of 'field 2' is first moved up and is movement interpolated with line 2 of 'field 1' moved down; it is then re-used so that line 2 of 'field 2' moved down is movement interpolated with line 3 of 'field 1' moved up. Subsequently, the information from the two field stores is read differently: for example, line 4 of 'field 1' is read at the same time as line 3 of 'field 2'. The process continues in this way until the information in the second spatial filter is recirculated: for example, line 5 'field 1', when it reverts to the original form.

It will thus be appreciated that while the 525-to-625 mode of conversion requires more complex interpolation than the 625-to-525 mode, the process is equally under the control of the designer and carried out with calculable accuracy.

References

1. R W Schafer and L R Rabiner, 'A Digital Signal Processing Approach to Interpolation', *Proc. IEEE,* Vol. 61, No. 6, June 1973.

2. J O Drewery, J R Chew and G M Le Couteur, 'Interpolation in Digital Line-Store Standards Conversion: A Theoretical Study', *BBC Engineering,* March 1973.

Digital Coding and Blanking

by A Bellis and P R Carmen

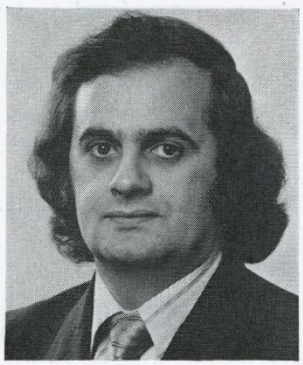

ARNOLD BELLIS,
BSc (Hons), graduated in electronic engineering from University College of North Wales and spent two years with the BBC. He then worked on the design of closed-circuit television with STC and subsequently joined Rank Cintel as project leader for the development of a new photoconductive telecine. He joined the IBA's Video and Colour Section in 1974. A Welsh-speaking Welshman, he now lives with his wife and son in Hampshire.

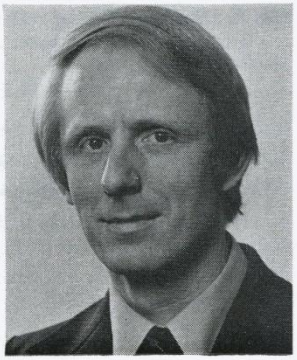

PETER CARMEN
completed a six-year apprenticeship at the Royal Radar Establishment, Great Malvern and subsequently worked on industrial television and military equipment with The Marconi Company. He joined the IBA's Experimental and Development Department in December 1973 and his work included the design of a digital colour coding system for DICE.

Synopsis

Digital coding and decoding are two of the many video processing requirements for DICE. Any analogue television signal can be converted into pulse-code-modulation (pcm) form by sampling; sometimes a multiple of the colour subcarrier frequency is used as the sampling frequency. Each sample constitutes proportions of the baseband signals determined by the relative positioning of the samples with the burst reference phase. Choice of the sampling frequency and the relative phasing is affected by practical considerations.

The derivation of baseband signals from the sample codewords and the construction of a coded signal from baseband signals both require consideration of the interaction between codeword groups in order to obtain acceptable frequency-response characteristics. In decoding, this is a function of chrominance separation from luminance in a spatial filter but, in coding, the response is dependent upon the way that the codewords are constructed.

Physical realisation of a practical coding system involves the use of multiple paths of adders, latches and switches appropriately controlled to combine percentages of the baseband components so that three codewords are produced for each cycle of subcarrier. A burst signal is constructed using the coding system operating on fixed baseband components injected during the blanking period.

Digital signals, when gated by a blanking waveform, have risetimes which are far too short for a television system; therefore it is necessary to modify the resultant blanking edges to reduce the rate of rise and fall so as to be within system specifications. By multiplying the digital signal by a time-varying multiplier over a suitable period at each blanking edge, the required risetime specification can be met. Simple multiplier coefficients can give a good approximation to the required curve. The mechanism for forming the blanking edges is used also to shape the baseband I and Q components in such a way that in subsequent coding the correct burst envelope is achieved, the multipliers now acting on gated-in static I and Q burst components.

The coding principle can be extended to produce static test signals including colour bars. Codewords scaled to the particular phase and amplitude of each bar are stored in memories and read out sequentially by a system of counters defining the duration of each bar and each transition.

The conception of a standards converter for two-way operation between the 525-line NTSC and 625-line PAL systems brought about a need to study many aspects of video processing in digital terms. The first DICE operated in the 525-to-625 direction only and one requirement was to decode the pcm digital NTSC signal into luminance and chrominance components. In this 525-to-625 direction, the main storage block was designed to handle the pcm NTSC signal directly from the analogue-to-digital converter (adc). This was convenient in that it provides a more economical form of storage in terms of cost and size than would the alternative of storing luminance and colour-difference signals. Re-arrangement of functional blocks when operating in the 625-to-525 direction therefore called for the use of a digital coder to convert baseband luminance and chrominance signals in Y, I and Q form into a pcm pseudo-NTSC waveform prior to the main storage unit. A second coder was required to provide the final 525-line NTSC-coded output.

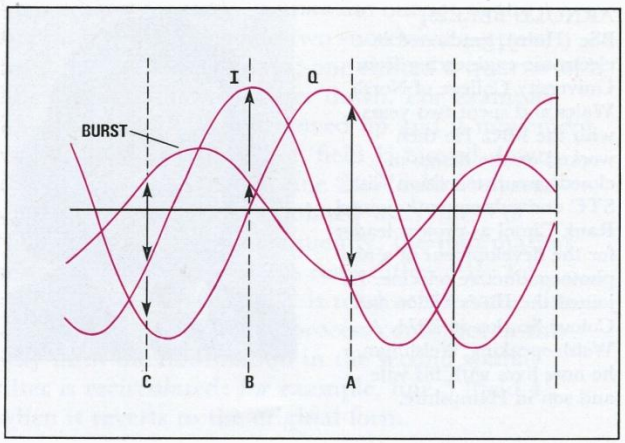

Fig. 1. Diagram showing relative phase of modulation on the I and Q axes relative to burst and to approximate sample phasors.

Decoding a Colour Signal

The functions of coding and decoding in DICE are the translation of baseband luminance and chrominance information into a phase- and amplitude-modulated subcarrier form of coded signal, and vice versa.

In this converter composite-coded signals are sampled at three times the NTSC subcarrier frequency ($3f_{sc}$) and it is convenient to use this same frequency for sampling baseband luminance information; this sampling rate is the lowest which can represent any phase and amplitude of subcarrier if the samples are equally spaced in time.

Figure 1 shows the way in which the specific cases of modulation on the I and Q axes and the burst are related by signed and scaled component phasors, A, B and C. If these phasors are transferred to the $R-Y$, $B-Y$ vector diagram, they appear as in Fig. 2, that is to say the A phasor lies precisely on the $B-Y$ axis (the burst being 180° to $+E_B-E_Y$) and the B phasor lies close to the I axis. The B phasor is along the 120° line to the $B-Y$ axis and therefore 3° lagging the true I axis at 123°. The C component is spaced 120° from both A and B. Thus, by referring again to Fig. 1, it can be seen that the burst is resolved by the A phasor lying at the negative peak and the B and C phasors have equal positive values at the 50% peak amplitude points. For reasons of practical convenience, these phasor positions have been adopted in DICE and the 3° error created by taking the B phasor to represent directly the I component of any modulated signal has been corrected by means described in other sections. The significance of adopting this approach becomes

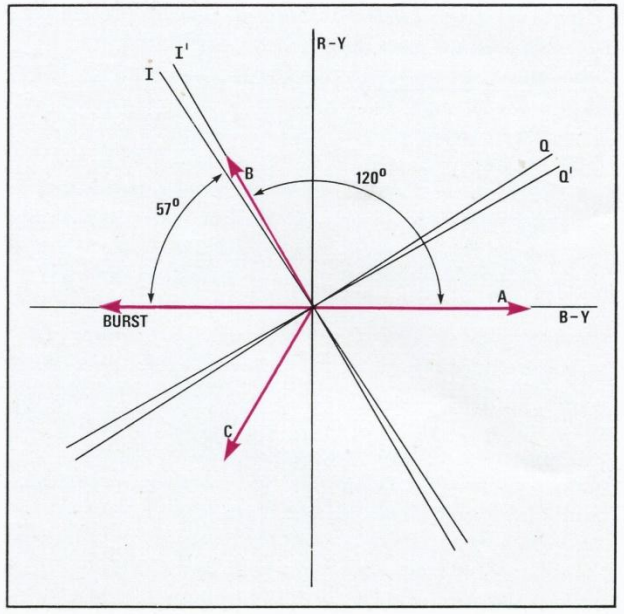

Fig. 2. Diagram showing phasor positions relative to I and Q axes and burst position.

apparent when the phasors are related to components of I and Q.

Consideration of Fig. 2 in terms of the Q and I contributions to the phasor magnitudes yields the following:

$$A = Q\cos 33° - I\cos 57° \quad \text{———(1)}$$
$$B = Q\cos 87° + I\cos 3° \quad \text{———(2)}$$
$$C = -(Q\cos 27° + I\cos 57°) \quad \text{———(3)}$$

The scaling factors introduced in the above equations are not compatible with the simple arithmetic functions which can be performed in a digital system, but if the approximation that equation (2) becomes B = I, then:

$$A = Q\cos 30° - I\cos 60° = (\sqrt{3}Q - I)/2 \quad \text{—(4)}$$
$$B = I \quad \text{———(5)}$$
$$C = -[Q\cos 30° + I\cos 60°] = -(\sqrt{3}Q + I)/2 \quad (6)$$

The equations derived above refer only to the chrominance content of the signal; the complete expression for the fully coded signal becomes:

$$A' = Y + (\sqrt{3}Q - I)/2 \quad \text{———(7)}$$
$$B' = Y + I \quad \text{———(8)}$$
$$C' = Y - (\sqrt{3}Q + I)/2 \quad \text{———(9)}$$

The only further consideration affecting the sample expressions is the reversal of subcarrier phase on alternate lines in the picture and this is accommodated by changing the sign of the chrominance expression on alternate lines.

Equations (7), (8) and (9) can be re-arranged to give the basic decoding expressions:

$$Q = (A' - C')/\sqrt{3}$$
$$I = [2B' - (A' + C')]/\sqrt{3}$$
$$Y = [A' + B' + C']/\sqrt{3}$$

where the magnitudes of I, Q and Y relate to A', B' or C' as shown in Fig. 3.

If decoding is performed using the three code words A', B' and C', the frequency responses obtained for I, Q and Y are unsatisfactory.

The chrominance characteristics are asymmetric about the subcarrier frequency, resulting in excessive quadrature crosstalk, and the luminance characteristic offers poor rejection to chrominance information resulting in crosstalk between the chrominance and luminance channels.

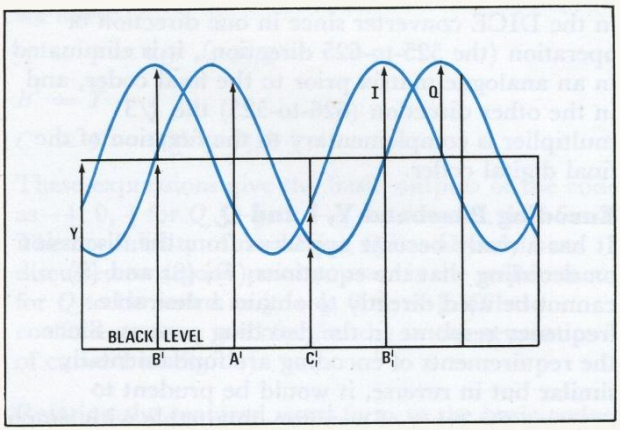

Fig. 3. Relationship between I, Q and Y components of a digital NTSC signal and the three-times subcarrier frequency samples.

A significant improvement can be obtained by incorporating codewords from the previous and following cycles of subcarrier. This produces chrominance characteristics which are virtually symmetrical about the subcarrier frequency.

Further improvement can be obtained if comb filtering is used to separate the chrominance and luminance signals before decoding is executed.[1]

The Decoding Function

Since the chrominance has been separated from the luminance, the decoding expressions are:

$Q=1\ (A_m - C_m)/\sqrt{3}$ and $I=B_m$ from (4), (5) and (6).

where A_m, B_m or C_m and codewords modified in the chrominance filter by the appropriate matrix coefficients.

The basis of a possible chrominance decoder is shown in Fig. 4. All codewords pass through delay '1', which equalizes the delay between the I and Q channels, to a latch which is clocked by input x to correspond with the B_m word time. Thus the clock rate is at $1/f_{sc}$. The output I therefore consists of B_m, B_{m+1}, etc. Delay '2' passes all codewords to an adder such that the delayed A_m word occurs at the same time as the undelayed negated C_m word. When the adder output is $(A_m - C_m)$ the latch after the adder is clocked by input y (at $1/f_{sc}$ rate) to give an output of $(A_m - C_m)$, $(A_{m+1} - C_{m+1})$, etc., or $\sqrt{3}Q$. The $\sqrt{3}$ does not present a problem

in the DICE converter since in one direction of operation (the 525-to-625 direction), it is eliminated in an analogue matrix prior to the final coder, and in the other direction (626-to-525) the $\sqrt{3}$ multiplier is complementary to the function of the final digital coder.

Encoding Baseband Y, I and Q

It has already become apparent from the discussion on decoding that the equations (7), (8) and (9) cannot be used directly to obtain a desirable frequency response in the decoding process. Since the requirements of encoding are fundamentally similar but in reverse, it would be prudent to examine the types of response obtainable with simple coding by direct use of (7), (8) and (9). Neglecting the luminance component which may be regarded as additive it can be seen that the Q output of a simple coder would be $\sqrt{3}Q/2$ for A, 0 for B and $-\sqrt{3}Q/2$ for C or for simplicity, gQ, 0, $-gQ$. The frequency response obtained from these points can be obtained from the general expression that $A(\omega)=a_0+2_j\ (a_1 \sin \omega T+a_2 \sin 2\omega T \ldots .+ a_n \sin n\omega T)$ in which $A\ (\omega)$ = amplitude as a function of frequency, a_0=zero, $a_1=g$ and $a_2=a_n=0$. T=time between samples. The plot obtained is curve 1 of Fig. 5 where the amplitude function has been normalised to unity at

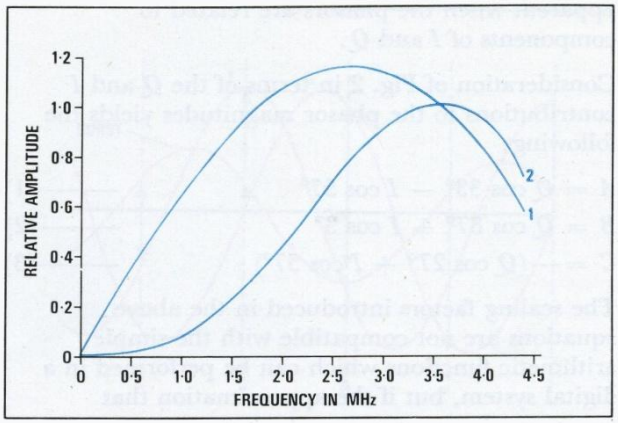

Fig. 5. (1) Plot of Q response of simple coder; (2) Plot of a modified Q response coder.

subcarrier frequency. Examination of this curve shows that at low frequencies, eg, 1 MHz, the Q output is high so that chrominance into luminance crosstalk would be excessive. (This could result in Hanover bars being observed on the coded picture.)

Consider again the general expression for the frequency response of the Q path: $A(\omega) = a_0 + 2_j\ (a_1 \sin \omega T + a_2 \sin 2\omega T + \ldots .a_n \sin n\omega T)$. It is

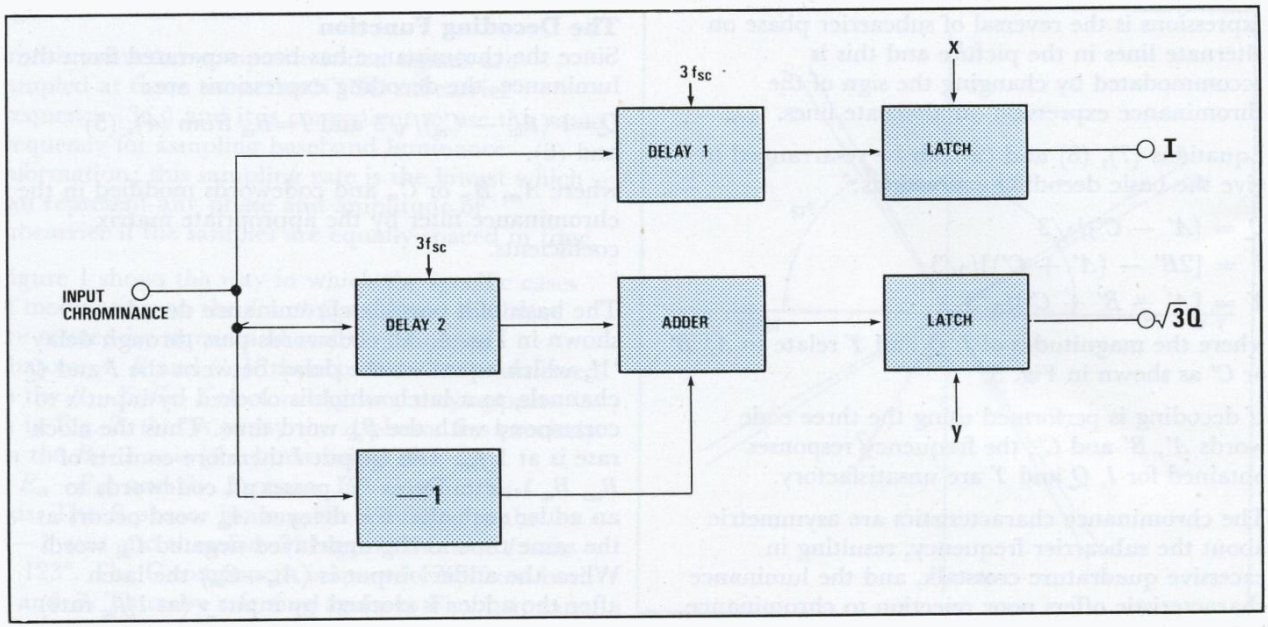

Fig. 4. Basic principle of chrominance decoder showing how I and $\sqrt{3}Q$ can be derived by delays, simple arithmetic and appropriately timed clocking by inputs x and y.

apparent that $A(\omega)$ can be modified by allowing terms after a_1 to have a value other than zero. Experience has shown that extending the codewords from three to five produces satisfactory results.

The Q output of the coder can be rewritten as:—

$-fQ, gQ, 0, -gQ, fQ$

and when $f=\frac{1}{2}g$ the modified response as shown in curve 2 of Fig. 5 is obtained. This offers a vast improvement over the original response and reduces the Q to luminance crosstalk at 1MHz by nearly 20dB. It also results in the centre of the passband being located at the subcarrier frequency.

Modifications to the construction of the coded words, whether applied to Q or I, do have side effects. Returning to the modified expression for Q output, it can be seen that for each complete f_{sc} period, the Q word is modified by the following and preceding Q words.

we have:

$A' = Y + (Q' - I)/2$

$B' = Y + I$

$C' = Y - (Q' + I)/2$

These expressions give the basic outputs of the coder as $-\frac{1}{2}$, 0, $\frac{1}{2}$ for Q, $-\frac{1}{2}$, 1, $-\frac{1}{2}$ for I and 1, 1, 1 for Y. The desirability of modifying Q has already been discussed so a good point to start is the requirement for Q to be coded as $\frac{1}{4}$, $-\frac{1}{2}$, 0, $\frac{1}{2}$, $-\frac{1}{4}$. (These coefficients were used for f and g in the generation of curve 2 in Fig. 5.)

Relating the required word form to the basic coder expression as shown below simple division results in the coefficients $\frac{1}{2}$, 1, X, 1, $\frac{1}{2}$. (The value of X is irrelevant since it is ignored by later stages in the coder.) The modified words are overlapping as mentioned previously and result in the Q gain of 3/2.

Word R	$-fQ_R$	gQ_R	0	$-gQ_R$	fQ_R						
Word S				$-fQ_S$	gQ_S	0	$-gQ_S$	fQ_S			
Word T							$-fQ_T$	gQ_T	0	$-gQ_T$	fQ_T
Effective Output				$-(f+g)Q$	$(f+g)Q$	0	$-(f+g)Q$	$(f+g)Q$			

The net effect of this modification is that the Q signal amplitude is changed by a factor $(f + g)$ and therefore must be allowed for prior to the coding function taking place. Another effect produced, as can be seen from curve 2 of Fig. 5 or the overlap of Q words, is a small reduction in Q bandwidth, although this is within acceptable limits.

Repeated Basic Q Codeword	$\frac{1}{2}$	$-\frac{1}{2}$	0	$\frac{1}{2}$	$-\frac{1}{2}$	0	$\frac{1}{2}$
Modified Q Codeword	$\frac{1}{4}$	$-\frac{1}{2}$	0	$\frac{1}{2}$	$-\frac{1}{4}$		
Modifying Coefficients	$\frac{1}{2}$	1	X	1	$\frac{1}{2}$		

Considerations for a Practical Coder

There are several different approaches to the design of any system, particularly in digital electronics. Sometimes it may be obvious which alternative offers the best solution in terms of performance, circuit complexity, physical size and cost; at other times, there may be several alternatives with apparently no particular points in favour of any one. The design of a coder suitable of meeting the criteria described earlier is a case in point; but one possible arrangement will be discussed in principle. Rewriting the equations relating I, Q and Y to codewords A', B' or C', and replacing $\sqrt{3}Q$ by Q'

Realisation of this part of the coder could therefore appear as in the block diagram in Fig. 6.

A single $1/f_{sc}$ duration Q word enters via S1 at the appropriate instant and is immediately halved, added with zero (assuming all previous Q words were zero), and passes on to the remaining coder functions. After a delay of one-third of the f_{sc} period, or the time of one sample codeword, the Q word appears at the Q input to the adder and is added to zero at the A input since S1 changes over

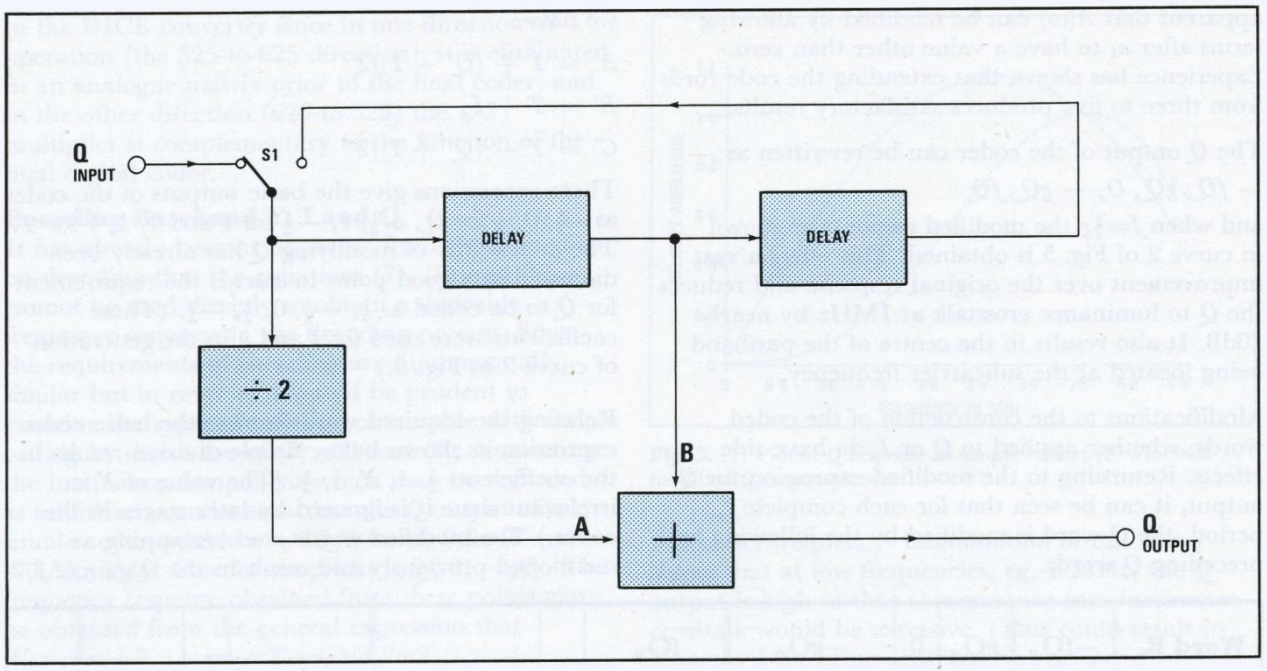

Fig. 6. Example of arrangement for producing the coefficients $\frac{1}{2}$, 1, X, 1, $\frac{1}{2}$ as applied to a single $1/f_{sc}$ duration Q word.

simultaneously. This zero is the fedback output of the second delay via S1. A further delay of one-third f_{sc} period reverses the situation so that B input is zero and A is one-half of Q. Thus the process is that B = the Q input, A = one-half Q input at the beginning of each period of $1/f_{sc}$ and A = one-half delayed previous Q word for the remainder of each period.

Hence, if $Q_o = 0$, $Q_1 = 1$, $Q_2 = 0$

the output is 0, 0, $\frac{1}{2}$, 1, $\frac{1}{2}$, 1, $\frac{1}{2}$, 1, 0, 0

or if $Q_o = Q_1 = Q_2 = 1$

is $\frac{1}{2}$, 1, $\frac{1}{2}$, $1\frac{1}{2}$, $1\frac{1}{2}$, $1\frac{1}{2}$, $1\frac{1}{2}$, $1\frac{1}{2}$ etc.

The modified Q signal derived from the above configuration can now be processed by a basic coder operating on the $-\frac{1}{2}$, 0, $\frac{1}{2}$ algorithm.

The block schematic diagram of Fig. 7 shows a possible way of generating the basic coder expressions. The Q signal is permanently divided by two and is switched by S1 to one side of the add/subtract stage. The I signal is fed to the other side of the add/subtract stage via S2 which selects I directly or $I/2$. Control to the add/subtract is arranged to give two additions and one subtraction for each subcarrier ($1/f_{sc}$) period. Thus the 'C_H' output of the 'adder 1' stage is $Q/2 + I/2$, I, $Q/2 - I/2$.

Control to 'adder 2' is arranged to give $-$, $+$, $+$ per subcarrier period, thus the output from the second adder is $Y - C_H$, $Y + C_H$, $Y + C_H$. S3 changes position at the start of each line so that the chrominance C_H is inverted line to line. Hence on the alternate lines the second adder output will be $Y + C_H$, $Y - C_H$, $Y - C_H$. This produces the chrominance phase inversion required on alternate lines in the NTSC system.

The detail design and construction of a coding system along the lines shown, or on any equivalent design must inevitably take account of the relatively high word rates used, i.e., most parts of the described system operate at $3f_{sc}$ (about 93 nS per word). Since all components have a finite delay which in some instances can be a significant percentage of the time of a word and that switches and adders require operating at precise instants of time, the use of fast logic such as Schottky ttl and carefully controlled clock-to-data timing is essential.

Another important consideration is the number of data bits required. In general, experience based on practical tests of digital coded signals shows that

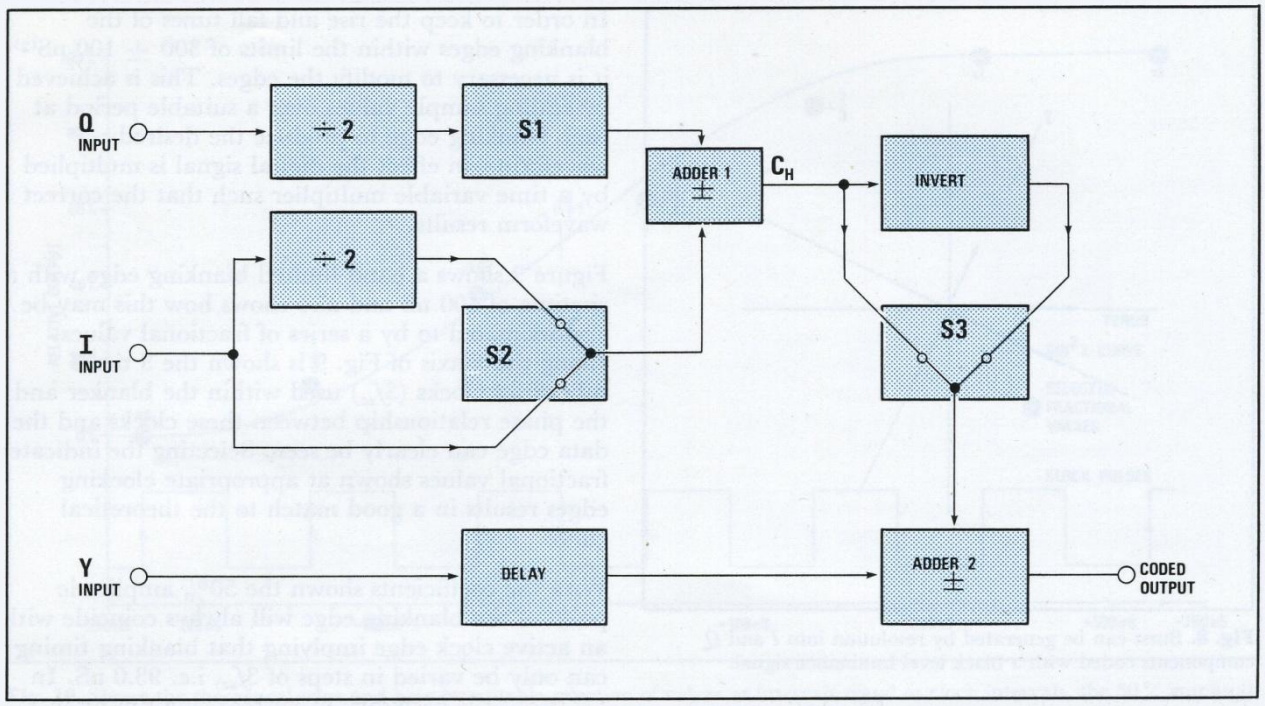

Fig. 7. Basic coder operation. During each $1/f_{sc}$ period S1 is on/off/on, giving $Q/2$, 0, $Q/2$; S2 gives 1/2, 1, 1/2, adder 1 gives $+$, $+$, $-$, so that $C_H = \frac{1}{2}Q + \frac{1}{2}I$, I, $\frac{1}{2}Q - \frac{1}{2}I$. The 'chrominance' signal C_H is added or subtracted to delay compensated luminance in adder 2. S3 is used to select the negated chrominance signal for the C codeword.

8 bits are required to represent each word without visibly contouring the picture. Thus, every adder, switch, delay or function incorporated in a data path has to accommodate 8-bit wide data.

Burst Generation and Correction of Phase

It was explained previously that sampling an encoded signal such that the A codeword is on the burst (B-Y) axis, resulted (assuming symmetrical sampling) in the B codeword lying at 3° to the I axis. The expressions derived for decoding and coding made use of the approximation for the sake of practical simplicity. Hence the components produced for I and Q when the coding system is used, as in DICE, to provide an NTSC output complete with burst, must be corrected. In the section on digital blanking, reference is made to the generation of burst (Fig. 8) by feeding baseband I and Q components, with appropriately shaped edges to the coding system. These components are scaled taking into account the 3° error approximation and result in a burst being generated such that the angular relationship between $+I$ and the burst is 57° thus eliminating the error. In addition, the component values are chosen to provide the correct level of burst relative to the video signal. The luminance input to the coding system during burst interval is video black level. The accuracy of the burst amplitude and phase is a function of the scaling between the analogue signal level finally required and the number of quanta levels used since the amplitude of burst is less than 25% of the total video signal. Thus the resolution of the digital burst is limited. However, with the correct choice of levels, the errors introduced can be reduced to insignificant proportions and can easily be contained within the NTSC specification limits.

Digital Blanking

If blanking edges with very short rise times were allowed in the television system, then the blanking of a digital video signal would be simplicity itself—by suitably gating the digital video signal with a blanking waveform the signal would be blanked but the resulting blanking edges would have rise and fall times equal to those of the logic family employed,

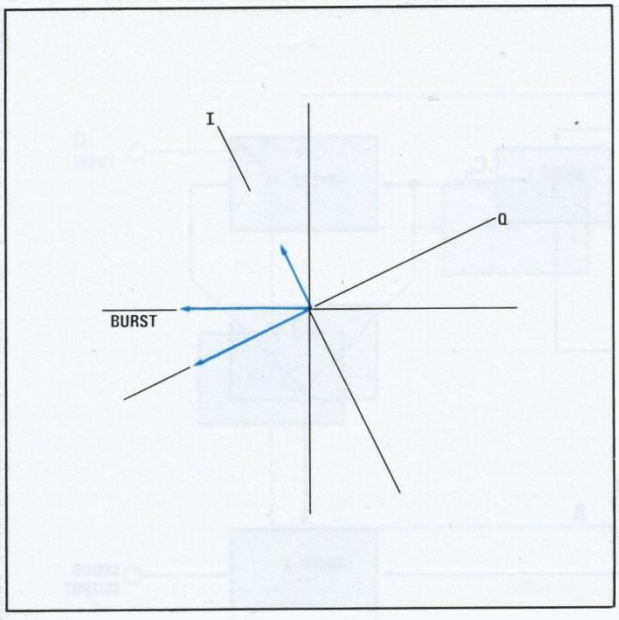

Fig. 8. Burst can be generated by resolution into I and Q components coded with a black level luminance signal.

and this would cause unacceptable overshoots on the blanking edge when the analogue signal had been filtered.

In order to keep the rise and fall times of the blanking edges within the limits of 300 $\pm$ 100 nS[2] it is necessary to modify the edges. This is achieved by scaling sample values over a suitable period at each blanking edge to produce the desired waveform—in effect the digital signal is multiplied by a time variable multiplier such that the correct waveform results.

Figure 9 shows a band limited blanking edge with a risetime of 300 nS and also shows how this may be approximated to by a series of fractional values. Along the x-axis of Fig. 9 is shown the 3 times subcarrier clocks ($3f_{sc}$) used within the blanker and the phase relationship between these clocks and the data edge can clearly be seen. Selecting the indicated fractional values shown at appropriate clocking edges results in a good match to the theoretical curve.

With the coefficients shown the 50% amplitude point of the blanking edge will always coincide with an active clock edge implying that blanking timing can only be varied in steps of $3f_{sc}$, i.e. 93.0 nS. In practice, it is necessary to set blanking timing in finer steps than 93.0 nS. In DICE steps of $6f_{sc}$, i.e., 46.5 nS have been chosen as timing increments.

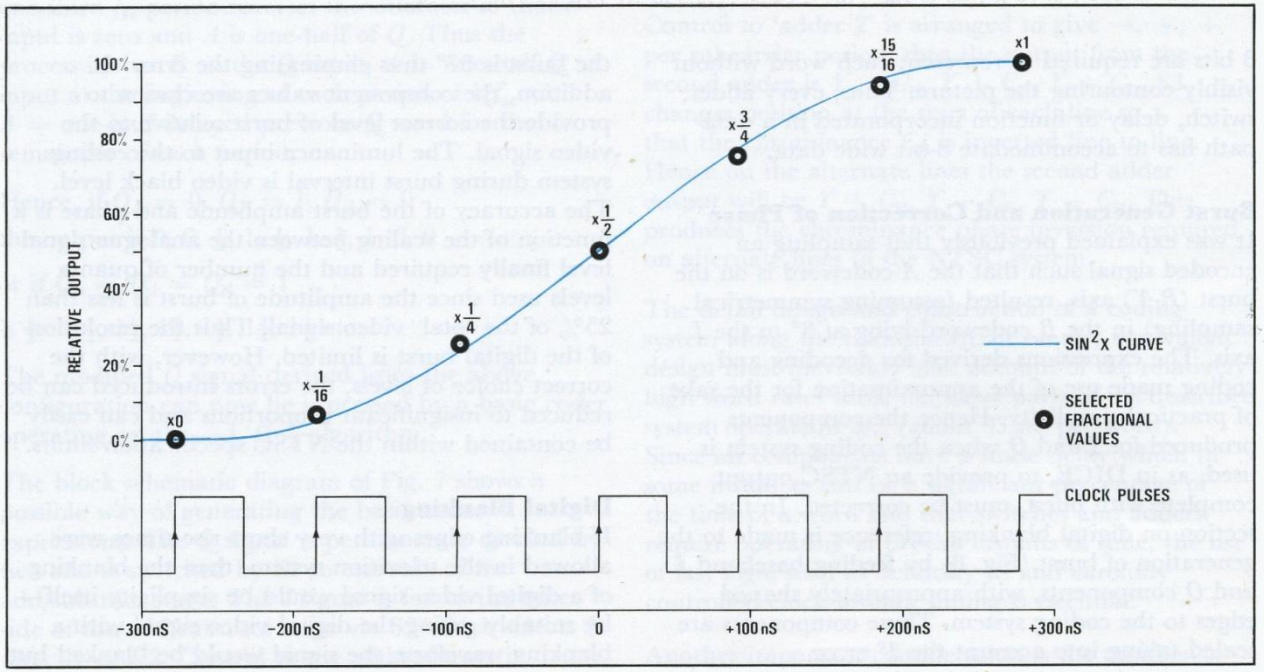

Fig. 9. Shows a theoretical $\sin^2 x$ blanking edge and how this can be approached by a series of fractional signal values at time intervals equal to the $3/f_{sc}$ clock pulse intervals.

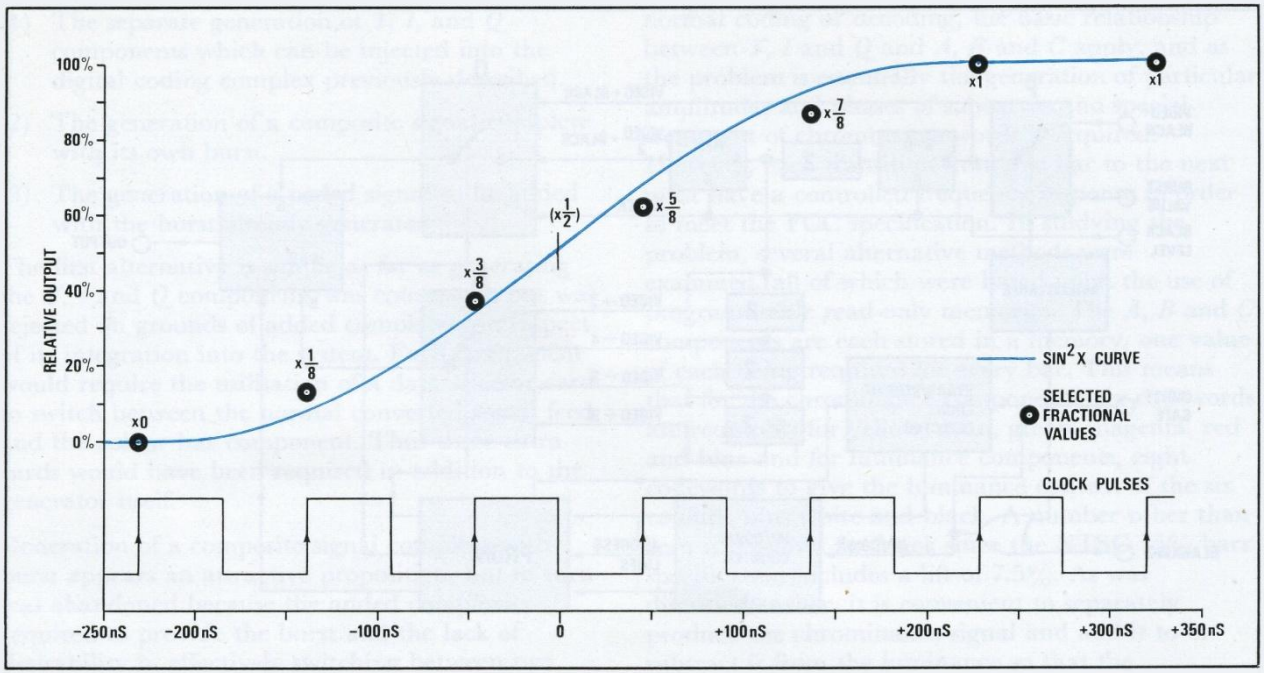

Fig. 10. Shows the theoretical edge and how by suitable selection of values at intervals equal to clock intervals, the 50% amplitude point of the curve may be set midway between clock active edges.

To move the 50% amplitude point of the blanking edge away from the clock active edge to a point midway between two active edges, a different set of multiplier coefficients are used. This new edge shape and multiplier coefficients are shown in Fig. 10 in relation to the $3f_{sc}$ clock pulses.

Since the coefficients chosen for this curve are not a symmetrical set, it is necessary to use yet another set of coefficients for a falling blanking edge with its 50% point midway between clock pulse active edges.

The fractional multiplier coefficients chosen can be simply synthesised from the full video signal amplitude, $\frac{1}{2}$ video, $\frac{1}{4}$ video, $\frac{1}{8}$ video, $\frac{1}{16}$ video by simple arithmetic operations.

For example $\frac{5}{8}$ video is generated as $\frac{1}{2}$ video + $\frac{1}{8}$ video, etc.

The basic fractional values required for synthesis can be generated quite simply by successively moving the significance of the digital video signal one bit to the right.

From the above considerations, the hardware requirements to implement a blanking system are: a *sequence controller*, a *coefficient generator* and an *arithmetic unit*. The basic system is shown in Fig. 11, which also contains other refinements, the significance of which will be explained later.

Control of the blanking sequence is by means of the up-down counter which counts 0 to 6 for a rising blanking edge and 6 to 0 for a falling blanking edge. The counter is started by the input edge and stopped by the start/stop circuit. The counter outputs feed a programmable read only memory (prom) which controls the selection of required fractional values for a given output value and also controls the output add/subtract arithmetic unit. One output of the prom is fed back to the counter via the start/stop of count detector in order to stop the counter once the correct state has been reached.

The video input to the blanker is in fact a video value superimposed on a dc value for picture black. It is obviously necessary to remove the black level value to give the true video value before further processing. This is carried out in the first subtractor. Similarly it is necessary to generate a value of $\frac{1}{2}$ video + black to which is added or subtracted one of the lesser fractions of video signal. This is generated in the first adder.

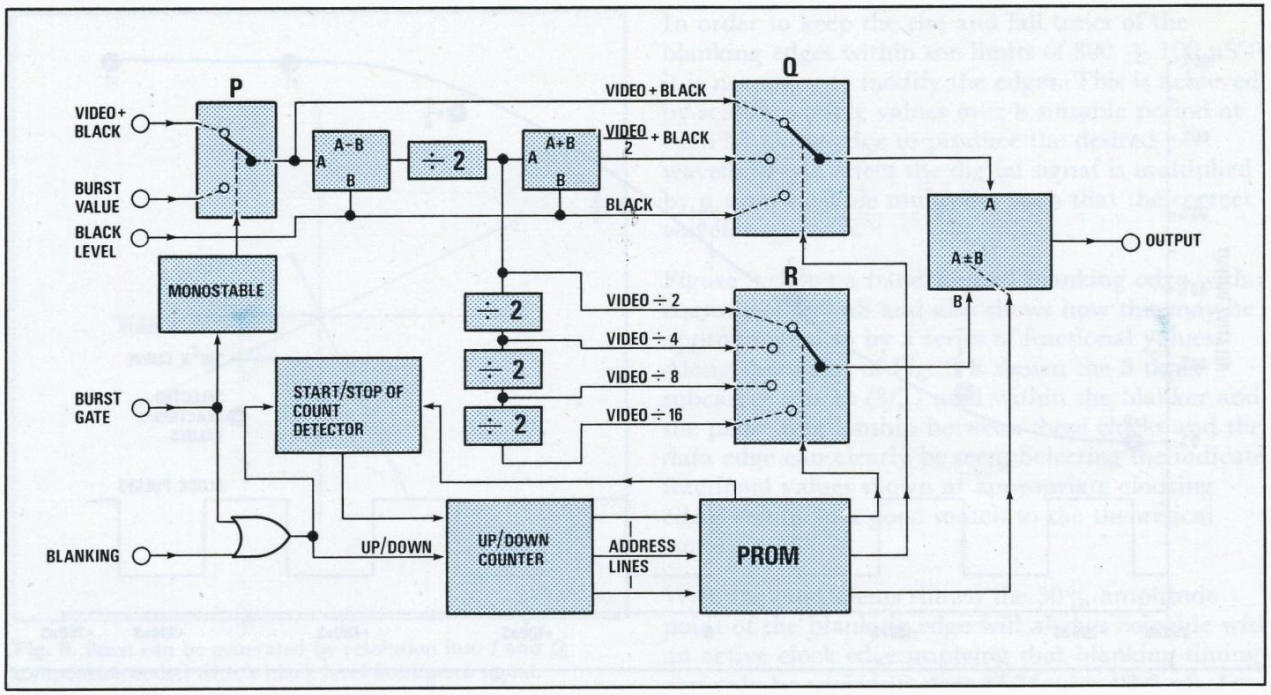

Fig. 11. Block schematic of the blanking system. The blanking edges are generated by forming fractional values of the video signal in a controlled manner over a given time period. The fractional values required are themselves generated from the values of 100% video, $\frac{1}{2}$ video, $\frac{1}{4}$ video, $\frac{1}{8}$ video and $\frac{1}{16}$ video by suitable arithmetic.

The *A* input side of the output adder/subtractor is fed with 'video + black', '$\frac{1}{2}$ video + black' or 'black' selected by data selector *Q*, whilst the *B* side of the adder/subtractor receives the lesser fractions of video via data selector *R*.

As explained earlier, three different sets of coefficients are required for the three possible forms of blanking edge. In the practical blankers there are therefore three prom devices, and not one as shown in Fig. 11. Each prom is programmed with the appropriate edge information. The correct prom is selected by taking account of whether the blanking edge is a rising or falling one and also comparing the timing of the incoming blanking signal with the $3f_{sc}$ clock pulses.

Burst Insertion

When DICE is converting in the 625-to-525 direction, the colour burst is inserted in the *I* and *Q* channel blankers as a baseband signal which is subsequently encoded by the digital coder. The burst envelope has its edges shaped by the same mechanism as the blanking edges are formed.

The required baseband *I* and *Q* burst values are gated into the blanking system by data selector *P* (Fig. 11) which is controlled by a monostable triggered by burst gate.

Switches on the blankers allow independent switching-off of burst or video, whilst a third switch allows the blanker to be switched off so that unblanked video is fed to the output for test purposes.

Generation of colour bars

One of the requirements of the DICE converter is that it should provide a standard colour-bar waveform output in either PAL or NTSC. When operating in the 525-to-625 conversion direction, PAL or SECAM 100% colour-bars are provided by the analogue coder, whereas in the 625-to-525 conversion direction, American 75% bars are generated using digital techniques.

There are three alternative methods open for the design of a digital colour bar system:

(1) The separate generation of Y, I, and Q components which can be injected into the digital coding complex previously described.

(2) The generation of a composite signal complete with its own burst.

(3) The generation of a coded signal to be added with the burst already generated.

The first alternative is simple as far as generating the Y, I and Q components was concerned, but was rejected on grounds of added complexity in respect of its integration into the system. Each component would require the utilisation of a data selector card to switch between the normal converted signal feed and the colour bar component. Thus three extra cards would have been required in addition to the generator itself.

Generation of a composite signal complete with burst appears an attractive proposition, but in turn was abandoned because the added complexity required to provide the burst and the lack of desirability in effectively switching between two sources.

The third alternative is a simplification of the second in that it still generates a coded signal directly but uses blanking period switching to obtain the burst provided by normal coding as described in the previous section. This approach is used in DICE.

Principle of Operation

The requirement is to produce codewords to represent each colour of the bar waveform. As for normal coding or decoding, the basic relationship between Y, I and Q and A, B and C apply, and as the problem is essentially the generation of particular amplitudes and phases of subcarrier, no special treatment of chrominance words is required. However, each transition from one bar to the next must have a controlled frequency response in order to meet the FCC specification. In studying the problem, several alternative methods were examined, all of which were based upon the use of programmable read-only memories. The A, B and C components are each stored in a memory, one value of each being required for every bar. This means that for the chrominance components, six codewords are required: for yellow, cyan, green, magenta, red and blue and for luminance components, eight codewords to give the luminance content of the six colours, plus white and black. A number other than zero is required for black since the NTSC 75% bars specification includes a lift of 7.5%. As was discussed earlier, it is convenient to separately produce the chrominance signal and add it to or subtract it from the luminance so that the subcarrier phase inversion from line to line can be realised. With this in mind, the basic form of a generator appears as in Fig. 12.

Chrominance and luminance memories are addressed in parallel by commands derived from sync and clock pulse inputs. The output of the chrominance memory is added to the luminance output on one line and subtracted on the alternate line by operating S1 at half line frequency. In principle, the plan looks

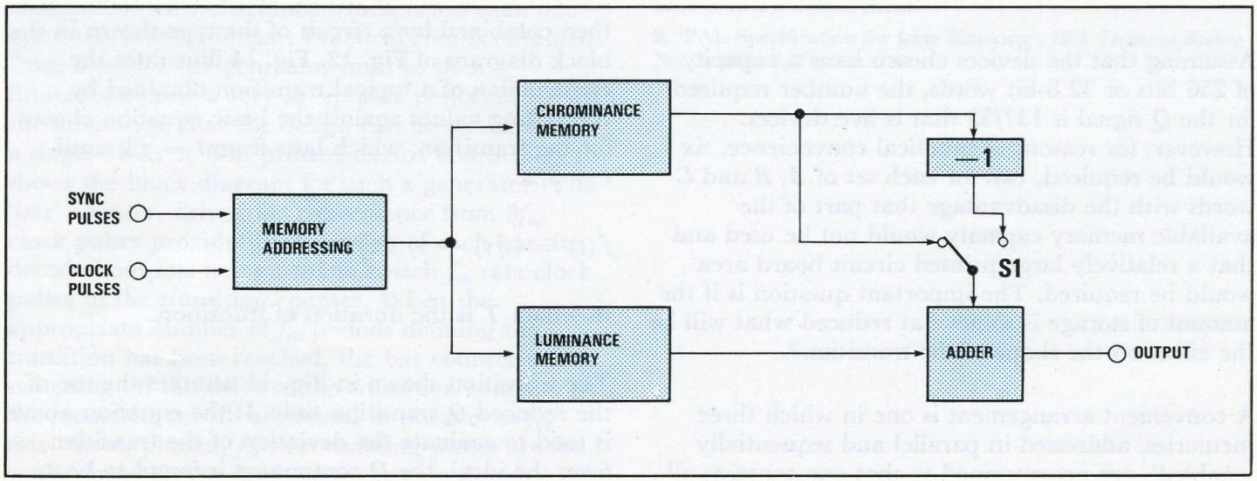

Fig. 12. Basic plan for colour bar generator.

straight-forward: some system of addressing employing counters, some suitably programmed memories and some easily implemented arithmetic. The problem centres mainly on the chrominance memory in establishing the number of codewords required to describe fully the bars and the transitions between the bars.

Interpretation of the FCC specification for the I and Q bandwidths indicates relative transition times of the order shown in Fig. 13. This means that, to accommodate the Q transition, six A, five B and six C codewords are required. The total number of codewords to represent the Q signal for line of bars is:

> Number of bars $\times$ number of codewords per cycle of subcarrier plus number of transitions $\times$ number of codewords per transition
> $= (6 \times 3) + (7 \times 17) = 137$

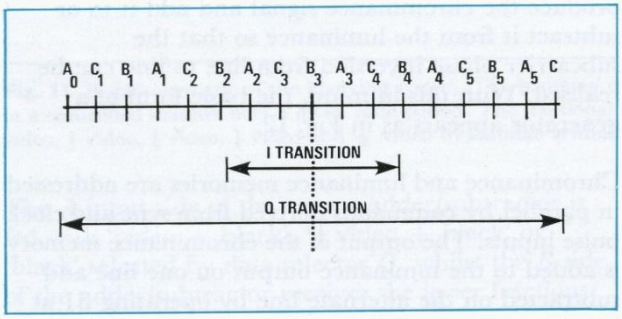

Fig. 13. Duration of I and Q transitions relative to codeword intervals.

Assuming that the devices chosen have a capacity of 256 bits or 32 8-bit words, the number required for the Q signal is 137/32 that is five devices. However, for reasons of practical convenience, six would be required, two for each set of A, B and C words with the disadvantage that part of the available memory capacity would not be used and that a relatively large printed circuit board area would be required. The important question is if the amount of storage is somewhat reduced what will be the effect on the shape of the transition?

A convenient arrangement is one in which three memories, addressed in parallel and sequentially 'enabled', are programmed so that one contains all A codewords, a second all B codewords and a third all C codewords. With 32 available locations to each, if 8 are used to provide the values for the bars (two for zero chrominance during white and black), then 24 remain to define transitions. If three locations are used to define interim points during a transition, it can be seen from Fig. 13 that assuming B_1 defines the last B codeword before the transition, B_2, B_3, B_4 give transition values and B_5 is the first B codeword of the next bar. This means that the transition time is limited to twelve codeword intervals, i.e., $12 \times 93.1\text{ns} = 1117\text{ns}$, instead of the ideal Q transition time of 18 intervals or 1676ns.

Another factor which requires consideration is the shape of the transition envelope. Apart from choosing a shape compatible with the specified bandwidth of the appropriate signal, it is desirable that the spectrum caused by the edge should not result in components at zero frequency or at minus the subcarrier frequency.

With this objective the frequency spectrum of the function used to define the edge is arranged to have zeros at multiples of the subcarrier frequency. Once a suitable expression for defining the transition has been established, values for the A, B and C components for both I and Q can be computed and added to give the total 'chrominance' components. These chrominance components are then programmed into a set of read-only memories. Similarly, the transition values for the luminance signal are computed and programmed into a further set of memories. Luminance and chrominance are then combined by a circuit of the type shown in the block diagram of Fig. 12. Fig. 14 illustrates the construction of a typical transition obtained by computing values against the basic equation chosen for the transition, which lasts from $t = \tau/2$ until $t = \tau/2$.

$$f(t) = (t/\tau) - 1/(2\pi) \sin 2\pi t/T$$

in which T is the duration of transition.

The transition shown in Fig. 14 assumes the use of the reduced Q transition time. If the equation above is used to evaluate the deviation of the transition from the ideal, the Q component is found to be within 1% of its final value at the defined duration

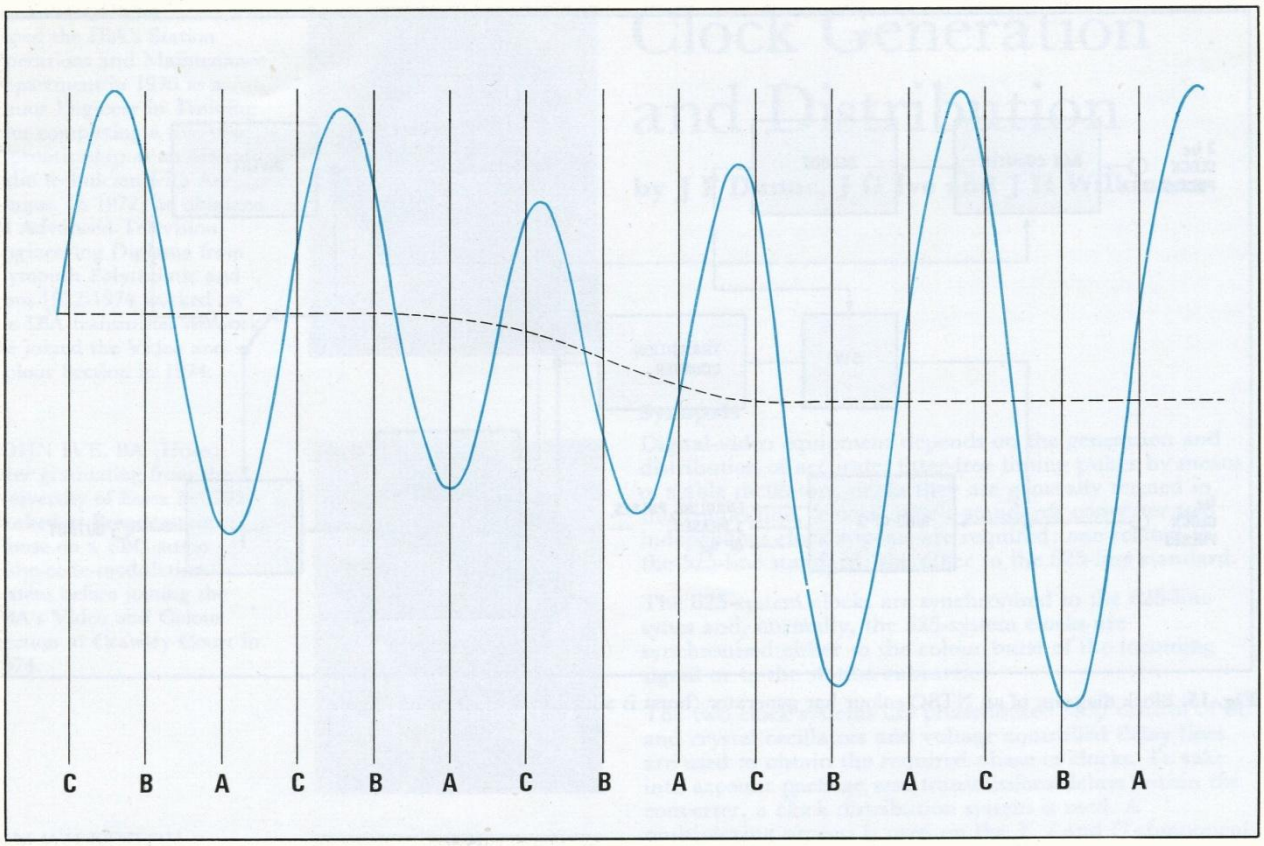

Fig. 14. Yellow-cyan transition constructed from computed values for codewords *A*, *B* and *C*. The luminance content is shown dotted.

of 1117ns compared with about 2.9% of its final value at 1117ns if the full period of 1676ns had been considered. This difference which is less than one quantum is insignificant. Practical results obtained from a colour bar generator built to these limitations show a very acceptable performance with the advantage that the design can be contained on a single 10-in × 6-in printed circuit board. Fig. 15 shows the block diagram for such a generator. The 'bar' counter, driven for convenience from $3f_{sc}$ clock pulses provides the duration of each bar, its decoded outputs being used to switch f_{sc} rate clock pulses to the transition counter. When the appropriate number of f_{sc} periods defining the transition has been reached, the bar counter starts, switching off the drive to the transition counter. As the memories are all addressed in parallel at f_{sc} rate, three, phased, f_{sc}-rate enabling pulses are fed to the memories to produce the required codeword at the right time.

References

1. J L E Baldwin and C Thirlwall, 'Spatial Filters' in this issue of *IBA Technical Review*

2. 'PAL Specification for Line Blanking', *IBA Technical Review* No. 2, page 4, Fig. 1.

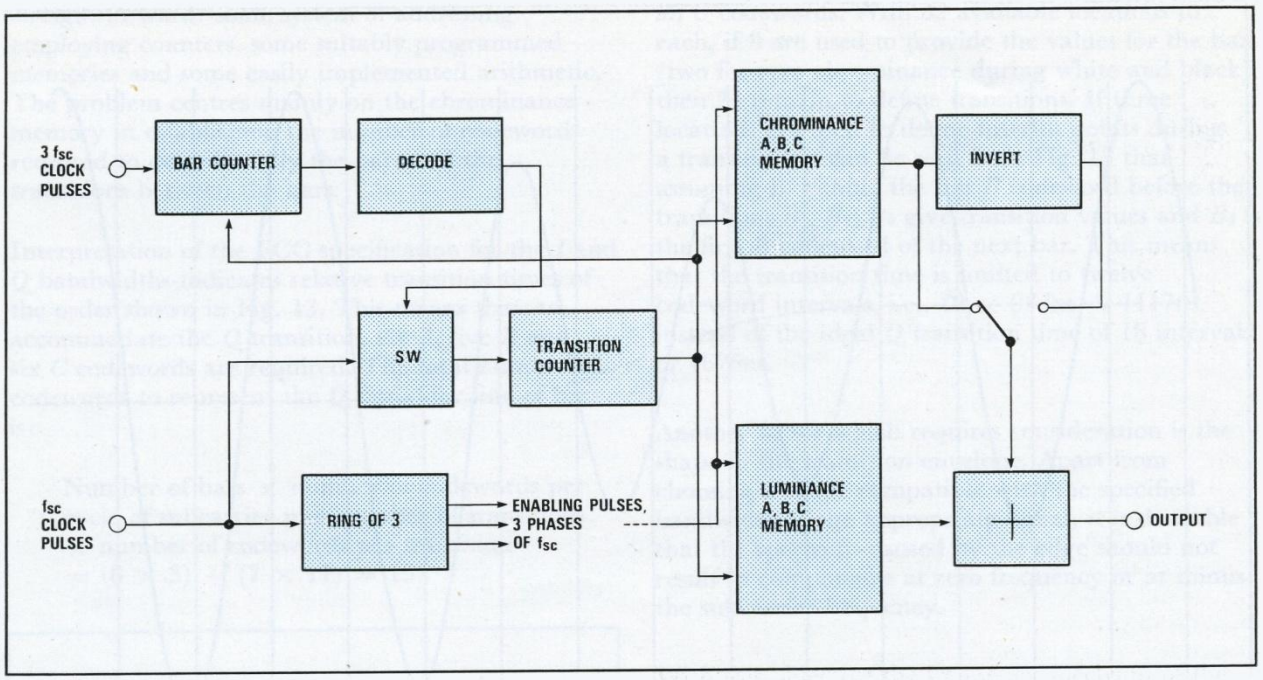

Fig. 15. Block diagram of an NTSC colour bar generator (burst is added in a subsequent stage).

Clock Generation and Distribution

by J F Dunne, J G Ive and J H Wilkinson

JOSEPH DUNNE joined the IBA's Station Operations and Maintenance Department in 1970 as a Junior Engineer in Training after completing a five-year apprenticeship as an aircraft radio technician with Aer Lingus. In 1972, he obtained an Advanced Television Engineering Diploma from Plymouth Polytechnic and from 1972-1974 worked on the IBA transmitter network. He joined the Video and Colour Section in 1974.

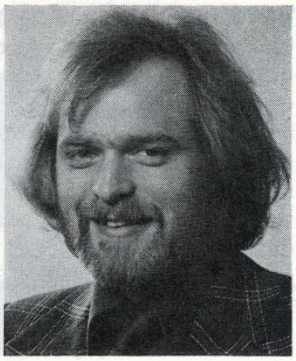

JOHN IVE, BA (Hons) after graduating from the University of Essex in 1972 worked at Broadcasting House on a BBC audio pulse-code-modulation system before joining the IBA's Video and Colour Section at Crawley Court in 1974.

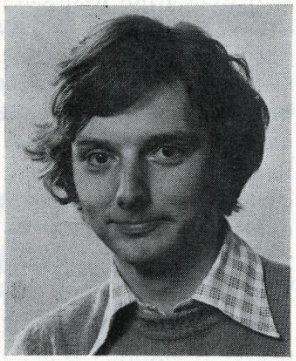

JIM WILKINSON studied at Sheffield Polytechnic after completing an apprenticeship with the Post Office. In 1972 he joined Marconi-Elliott Avionics, working on aircraft display systems at Rochester before joining the IBA's Video and Colour Section in 1974.

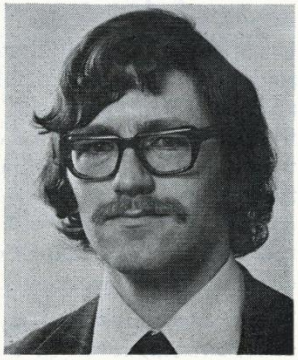

Synopsis

Digital-video equipment depends on the generation and distribution of accurate, jitter-free timing pulses by means of stable oscillators, or, as they are generally termed in this connection, 'clocks'. For a standards converter two independent clock systems are required: one related to the 525-line standard; the other to the 625-line standard.

The 625-system clocks are synchronised to the 625-line syncs and, normally, the 525-system clocks are synchronised either to the colour burst of the incoming signal or to the station subcarrier.

The two clock systems use phase-locked loop control of *LC* and crystal oscillators and voltage controlled delay lines are used to obtain the required phase of clocks. To take into account package and transmissions delays within the converter, a clock distribution system is used. A multiplexing process is used on the Y, I and Q component signals to minimise hardware requirements in the converter.

Two independent clock systems are required in DICE: clocks relating to the 525-line system, and those relating to the 625-line system. Because DICE converts 525-line composite-video to a digital form in the forward (525-to-625) direction, the 525 system clocks must be extremely stable and free of jitter. If there are any timing errors on the 525 clocks, then there will be hue errors on the analogue output when the converter is working in the forward direction.

The 525-line composite video input is sampled at three times subcarrier frequency ($3f_{sc}$) and the phase of these sampling clocks with respect to the $+I$

component of the chrominance signal is fixed. This enables the digital demodulation of the composite signal to be carried out simply; further, the 525-line clock system must be capable of operating from a monochrome source not having f_{sc} burst.

The 625-line system clocks do not need to be synchronised to the 625-line subcarrier frequency because these clocks are used only with baseband signals. The encoding and decoding of the 625-line system is carried out with analogue circuitry; hence the 625 clocks are synchronised to the 625-line syncs.

It follows that two clock-distribution systems are used in the converter; one is required to distribute the clock signals on the input side of the converter; the other distributes the clock signals on the output side. This implies that the distribution systems must, therefore, distribute a different set of clock timing pulses for each direction of operation of the converter.

Design philosophy

The function of the *525 master oscillator* is to generate clock timing signals synchronised to the 525-line NTSC system. In normal operating conditions, the clocks are synchronised to the colour burst of the composite video input in the forward direction; to station subcarrier in the reverse direction. The only exception is when the converter is operating in monochrome in the forward direction; in this case the clocks are locked to the 525-line syncs.

For colour, the clocks should be capable of varying their phase with respect to the colour subcarrier. In the reverse direction, the phase of the clocks determines the phase of the chrominance and burst on the output video signal. The phase of the subcarrier, therefore, must be variable over 360° so that the converter output can be mixed with other sources. In the forward direction (525-to-625), the 525 clocks perform the sampling of the incoming NTSC video signal in the analogue-to-digital converter (adc); hence the position and stability of these clocks are vital to the correct operation of the subsequent processing.

As noted earlier, it has been found that the hue on an NTSC incoming signal may be incorrect; to compensate for this, it has been arranged in DICE that the position of the 525 clocks can be varied with respect to the colour burst, hence adjusting the position of the sampling pulses in the adc. This effectively produces a hue change on the 625-line output picture. However, for normal colour operation a digital phase stabiliser is employed to correct any sampling errors. This stabiliser accepts the digitised video from the adc and, from the samples taken of the burst, it produces a control signal to compensate for any error in the position of these sampling pulses.

To generate system clocks from colour subcarrier, three phase-locked-loops have been employed. A phase-locked-loop produces an output frequency which is locked to an input frequency: see Fig. 1. One advantage of such a circuit arrangement is that the output frequency may be made more stable than the reference frequency or source in terms of both jitter and amplitude noise. Also, by inserting a frequency divider (dividing by n, where n is an integer) into the feedback circuit, the output may be at a frequency n times greater than the reference frequency. If a delay line is inserted in the feedback circuit, then the output will be 'advanced' by an equivalent amount.

In the forward direction (525-to-625), the 525 clocks must have less than 0.5 ns jitter, approximately equal to 0.65° of subcarrier. This degree of stability is achieved by using a crystal oscillator phase-locked to the subcarrier. However, since it is a requirement that DICE should be able to cope with 525-line monochrome sources, such sources may have a tolerance of $\pm 0.1\%$ in frequency stability of the line syncs, from which the clock frequencies must now be derived. Because of this wide tolerance, the oscillator must be an *LC* type rather than crystal-controlled. This results, of course, in clock signals of lower phase stability, but this is not of significance since there are no chrominance components in the signal.

The outputs must have a common coincident edge and this is achieved by using a single oscillator, followed by a divider chain. The oscillator must, in effect, be at the lowest common multiple of the various frequency components for which the clocks are required. These are: f_{sc}, $2f_{sc}$, $3f_{sc}$, $5f_{sc}$, $6f_{sc}$ (where f_{sc} is the colour subcarrier frequency which for NTSC is 3.579545 MHz $\pm$10 Hz). The lowest common multiple for these components is $30f_{sc}$ which is 107.386350 MHz $\pm$300 Hz. While it would be possible to make a crystal-controlled oscillator for this frequency, by using a suitable crystal in an overtone mode, it would

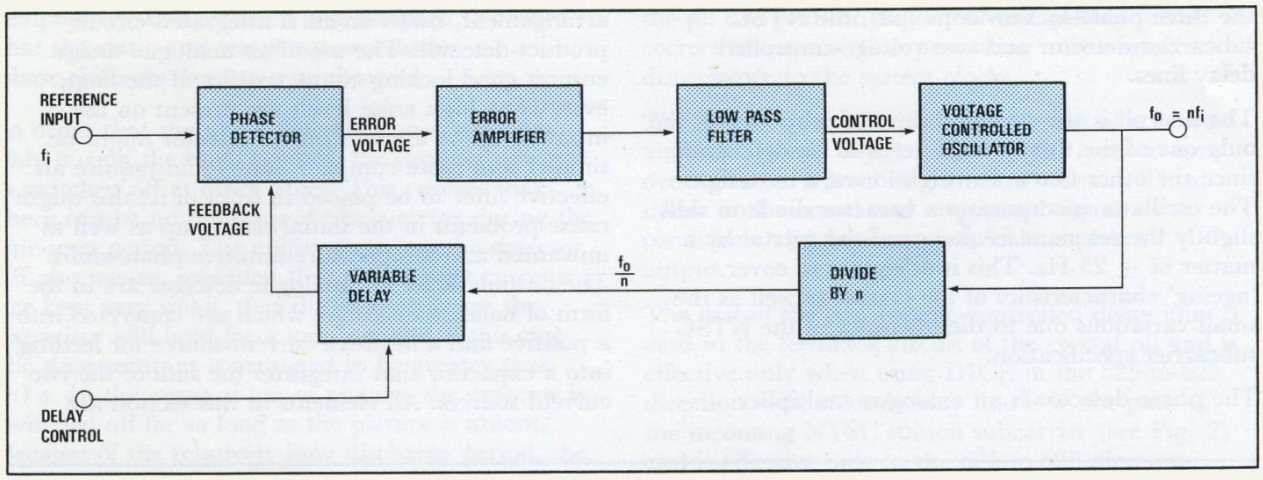

Fig. 1. A phase-locked-loop. The basic phase-locked-loop system has the ability to generate an output signal whose frequency (*fo*) is *n* times the input frequency (*fi*) and the phase of the output can be altered with respect to the input. Note that the variable delay network adds a delay in the feedback path; thus the output signal will be in advance of the input signal by the same amount.

be difficult to make on a production basis. So in practice an *LC* oscillator is used at $30f_{sc}$ which is phase-locked to a crystal oscillator at $2f_{sc}$ (7.159090 MHz). The phase-locked-loop in this case must have a fast response in order to minimise high-frequency phase-modulation of the *LC* oscillator.

Such an arrangement can be made to cope with a monochrome source in a straightforward manner: an *LC* oscillator at $2f_{sc}$ is used as the reference to lock the $30f_{sc}$ *LC* oscillator, while the output from the $2f_{sc}$ *LC* oscillator is divided by 455 to provide a signal at line rate so that the oscillator can be phase-locked to the line sync signals (there are 455/2 cycles of subcarrier in one line period of the NTSC system).

A detector is used to decide whether the incoming 525-line video signal is colour or monochrome: if colour, then the $30f_{sc}$ *LC* oscillator is phase-locked to the $2f_{sc}$ crystal oscillator; if monochrome, then the $30f_{sc}$ *LC* oscillator is phase-locked to the $2f_{sc}$ *LC* oscillator. It is possible, by means of a switch on the control panel, to override the detector during colour operation.

The crystal oscillator is always used in the reverse (625-to-525) direction: this is because the DICE specification states that station subcarrier will always be present. The $30f_{sc}$ *LC* oscillator is always locked to the crystal oscillator to ensure that the clocks are stable and suitable for generating a composite NTSC output signal.

It is a requirement that the phase of the output subcarrier (in both directions of operation) should be variable over 360° of subcarrier phase so that the output may be mixed with other sources. In the 525-to-625 direction, the 625 coder provides this facility; for 625-to-525 it is provided by the 525 master oscillator. In the 625-to-525 direction the subcarrier is routed through a voltage-controlled delay-line before it is used to phase-lock the clock oscillators (see Fig. 2). It is thus a simple matter to apply the correct voltage to this delay to set up the required subcarrier phase.

Because of the inherent phase inversion of subcarrier frequency, line by line (relative to the line syncs), it is necessary to generate line-coherent clocks for controlling some arithmetic operations. These clocks will then be in the same place on a line-by-line basis. Further, a subcarrier phase-identification signal (a square-wave at half line-frequency) is used to ensure that the data is correctly interpreted.

Details of Operation

The 525 master oscillator consists of four cards and one module: three of the cards consist of logic elements; the fourth, together with the module, forms the analogue parts of the system.

To consider first the analogue parts: these consist of

the three phase-locked-loops (pll), the NTSC subcarrier detector and two voltage-controlled delay lines.

The first pll is the crystal-controlled loop; this is the only one of the three which needs to be described since the other two are straightforward in design. The oscillator incorporates a varactor diode to shift slightly the resonant frequency of the crystal by a matter of $\pm$ 25 Hz. This is adequate to cover 'ageing' characteristics of the crystal as well as the small variations due to the tolerance of the NTSC subcarrier specification.

The phase detector is an analogue multiplier arrangement, based on an rf integrated-circuit product-detector. The use of an analogue design ensures good locking characteristics of the loop, even when high noise levels are present on the input signal. A digital form of detector might be simpler and more compact, but would require an effective filter to be placed in front of it; this might cause problems in the initial setting up as well as unwanted and temperature-sensitive phase-shifts. The outputs from the analogue detector are in the form of balanced voltages which are converted into a positive and a negative current-source for feeding into a capacitor that integrates the sum of the two current sources. All elements in this section are

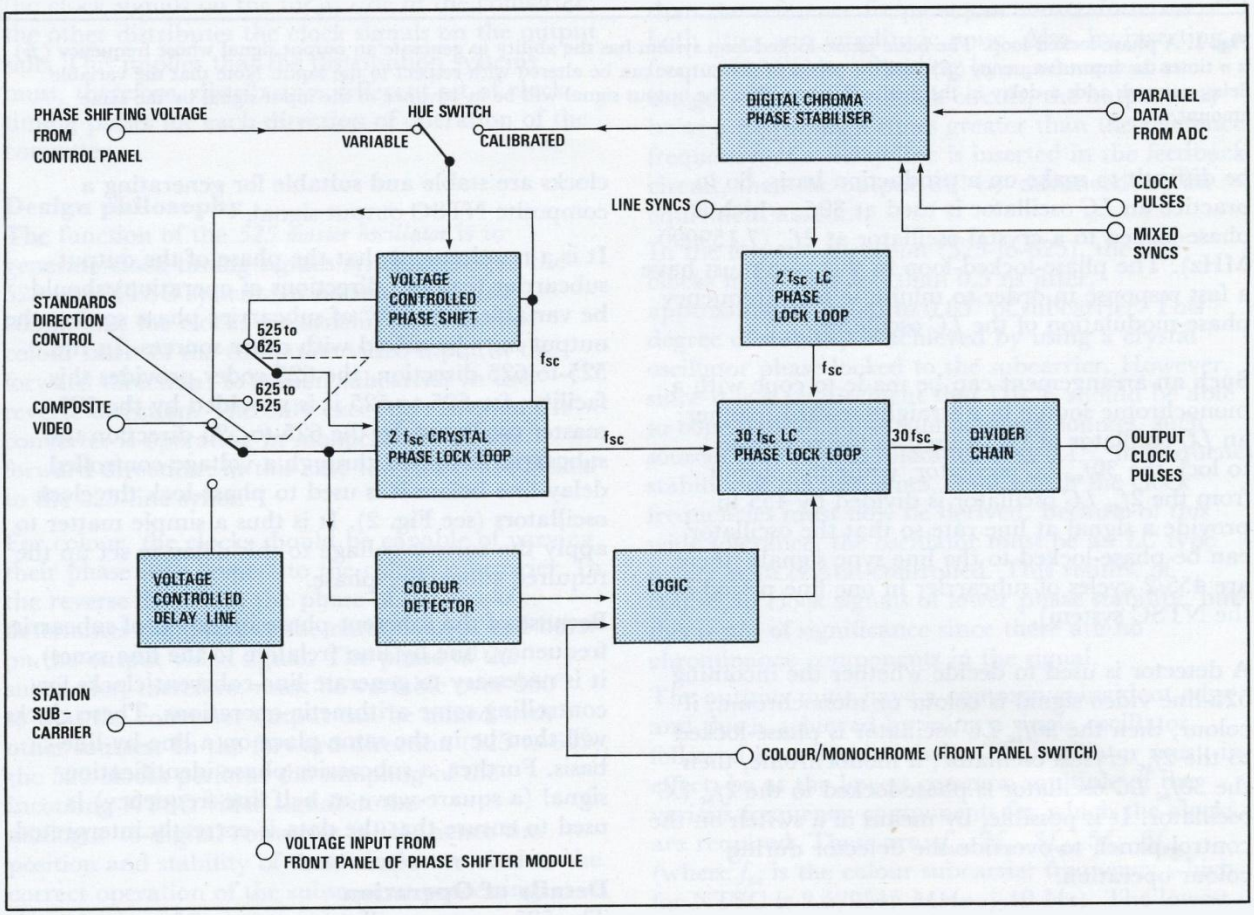

Fig. 2. Block schematic of the '525 master oscillator'. In the forward direction the $2f_{sc}$ crystal phase-locked-loop (p11) is used for colour sources and the $2f_{sc}$ *LC* p11 is used for monochrome sources. In the reverse direction the $2f_{sc}$ crystal p11 is always used. The voltage-controlled phase-shift network is used only in the forward direction. The amount of phase shifting is determined by either the chroma phase stabiliser (for calibrated phase) or by the adjustment on the control panel (for manually variable phase). This phase shift will be reflected through to the output clocks. Note that this phase-shifting facility is not used on monochrome sources where there is no colour subcarrier burst present.

temperature-compensated and designed to ensure that the static and temperature-sensitive phase-shifts are kept to a minimum.

In order that the phase detection process occurs only during the burst periods, the product detector is switched off at other times. This ensures that there can be no build up of phase-errors during the line-scan period. The ability to switch the detector off also means, provided that the leakage currents are kept very small, that the voltage across the capacitor will hold for a long period: in this case the time-constant is arranged to be greater than 0.1 s. In the event of loss of picture, the detector is switched off for so long as the picture is absent. Because of the relatively long discharge period, the oscillator is held at the correct frequency for a matter of several frames. Should the picture not return after 0.1 s, the oscillator frequency will drift slowly to a mean value which should be within 10 Hz of subcarrier frequency. When the signal finally returns, the error will be relatively small and the pll can pull the oscillator very quickly to the correct frequency, thus presenting only a small disturbance to the system clocks.

The NTSC subcarrier detector consists of a tuned amplifier (tuned to f_{sc}) whose output is set to turn over on a burst level 20 dB below normal. Thus, under both normal and attenuated signal conditions, the detector gives 'colour present' output.

The first of the two voltage-controlled delay lines is used in the feedback circuit of the crystal pll and is effective only when using DICE in the 525-to-625 direction. The second delay line is in the path of the incoming NTSC station subcarrier (see Fig. 2) and is effective only in the 625-to-525 direction.

The first delay line corrects for hue changes on NTSC pictures; it has a range of adjustment of $\pm$ 40° of subcarrier, and is controlled from either the digital phase stabiliser or from the hue adjusting switch on the control panel. When the 'hue' switch

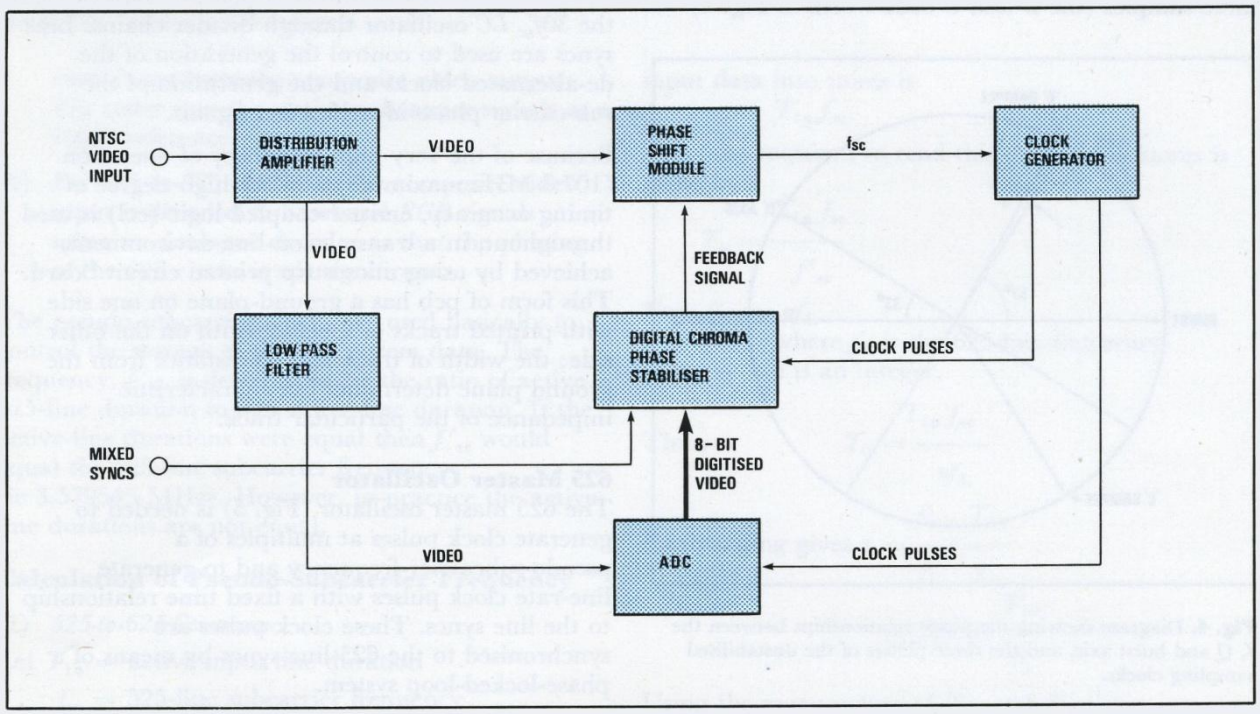

Fig. 3. Block schematic of chroma phase stabiliser system. The function of this system is to stabilise the sampling position of the analogue-digital converter (adc) clocks with respect to the burst phase on the incoming video. The chroma phase stabiliser accepts digital samples of the burst from the adc and from these it generates a feedback signal which is used to control the phase of the adc clocks. The chroma phase stabiliser has a built-in 3° offset to ensure that one of the samples is actually on the $+I$ axis, as shown in Fig. 4. Note that this system is used only in the 525-to-625 direction.

on the panel is in the 'cal' position, a voltage feedback from the digital phase stabiliser controls the phase adjustment. This feedback ensures a 3° phase error in the sampling position (see Fig. 3) of the NTSC video, so that one of the samples lies on the $+I$ axis and guarantees that the arithmetic operations in subsequent stages will be correct. If the picture hue is wrong, then the 'var' position of the 'hue' switch can be selected; in this case, the phase shifting voltage is controlled manually from the control panel.

The second delay-line consists of a standard two-element, voltage-controlled unit with a range of 180° of subcarrier phase; two hyper-abrupt junction varactor diodes are used to obtain this range of adjustment.

Digital Circuitry

The function of the digital phase-stabiliser is to ensure that the sampling positions in the adc module are correctly positioned. To achieve this two of the three samples (the B and C code words of Fig. 4)

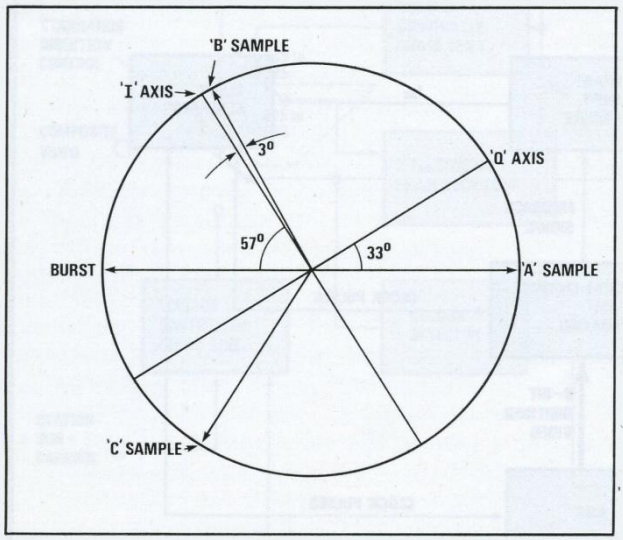

Fig. 4. Diagram showing the phase relationships between the I, Q and burst axis, and the three phases of the unstabilised sampling clocks.

which would normally be equal in the correct sampling position are digitally summed during each burst period; the accumulated error, representing phase error, is then used to generate a dc control voltage which is used to control the first delay line. This feedback system, it should be noted, is operational only when the 'hue' switch is in the 'cal' position.

The values of the two samples, B and C, are modified by a digital arithmetic process so that the system must give a 3° offset, thus ensuring that one of the samples is on the $+I$ axis. This is essential since the arithmetic operations in the NTSC digital demodulator section depend on one of the samples being on the $+I$ axis. This system of phase feedback, however, is not required in the 625-to-525 direction.

The function of the digital divider chain is to generate clock pulses at multiples of the subcarrier frequency (ie at f_{sc}, $2f_{sc}$, $3f_{sc}$, $5f_{sc}$ and $6f_{sc}$). All these clock pulses must have a common coincident edge. Owing to the inherent phase inversion of the NTSC subcarrier, line by line, it is necessary to generate further clocks having the same phase (with respect to the line sync edge) on a line-by-line basis. These additional clocks are known as *de-alternated clocks*.

All the 525 master oscillator clocks are derived from the $30f_{sc}$ LC oscillator through divider chains. Line syncs are used to control the generation of the de-alternated clocks and the generation of the sub-carrier phase-identification signals.

Because of the very high frequency of operation (107·3 MHz maximum) and the high degree of timing accuracy, emitter-coupled-logic (ecl) is used throughout, in a transmission-line environment, achieved by using microstrip printed circuit board. This form of pcb has a ground-plane on one side with printed tracks of specific width on the other side; the width of track and its distance from the ground plane determines the characteristic impedance of the particular track.

625 Master Oscillator

The 625 master oscillator (Fig. 5) is needed to generate clock pulses at multiples of a pseudo-subcarrier-frequency and to generate line-rate clock pulses with a fixed time relationship to the line syncs. These clock pulses are synchronised to the 625-line-syncs by means of a phase-locked-loop system.

It is unnecessary to use the 625-line system subcarrier as the reference frequency. This is because:

(1) For 525-to-625 operation an analogue coder is

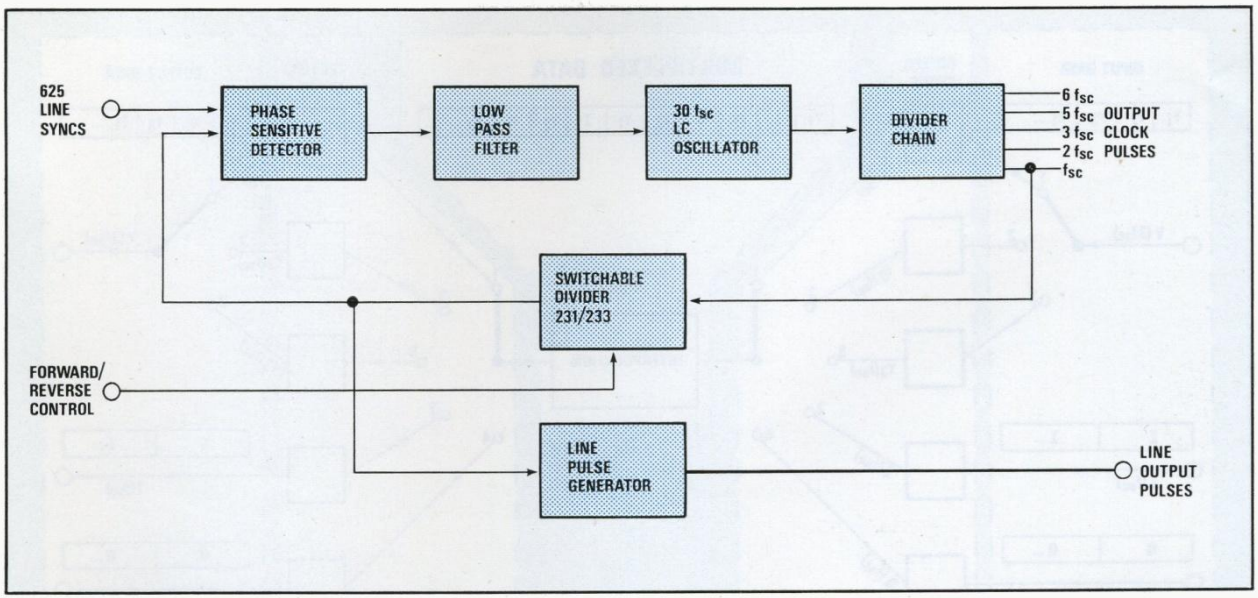

Fig. 5. Block schematic of the 625 master oscillator. The phase-locked-loop is used to control the frequency of the $30f'_{sc}$ oscillator and hence the frequency of the output clocks. The output frequency f'_{sc} is 231 times the 625-line frequency for 525-to-625 conversions and is 233 times the 625-line frequency for 625-to-525 conversions. The line output pulses can be independently set in any position relative to 625-line syncs.

used to generate the composite video output; this coder uses the station subcarrier signals as a stable reference.

(2) For 625-to-525 operation the composite video input is decoded to baseband *RGB* signals using an analogue decoder so that the phasing of the 625 clocks is not significant.

The pseudo-subcarrier clocks are used basically to control the storage of the 625 system data. The frequency, f'_{sc}, is determined by the ratio of active 625-line duration to 525 active-line duration. If the active-line durations were equal then f'_{sc} would equal the 525-line subcarrier frequency (ie 3.579545 MHz). However, in practice the active line durations are not equal.

Calculation of Pseudo-Subcarrier Frequency

(1) *525-to-625 Conversion:*

Let T_{in} = active input line duration

f_{sc} = 525-line subcarrier frequency
= 3.579545 MHz

T_o = Time required to read the input data at the output clock rate (f'_{sc})

The number of clock pulses required to read the input data into stores is

$$= T_{in} f_{sc}$$

The time required to read this data out of stores is given by

$$T_o = \frac{T_{in} f_{sc}}{f'_{sc}}$$

Now $f'_{sc} = nf_L$

where f_L is the 625-line-frequency
n is an integer.

Thus,
$$T_o = \frac{T_{in} f_{sc}}{nf_L}$$

Rearranging gives
$$n = \frac{f_{sc}}{f_L} \cdot \frac{T_{in}}{T_o}$$
$$= 229.09 \frac{T_{in}}{T_o}$$

Using the mean values of T_{in} and T_o gives

$$n = 229.09 \times \frac{52.606}{51.95} = 232$$

and $f'_{sc} = nf_L = 232 \times 15.625 \times 10^3 = 3.625$ MHz

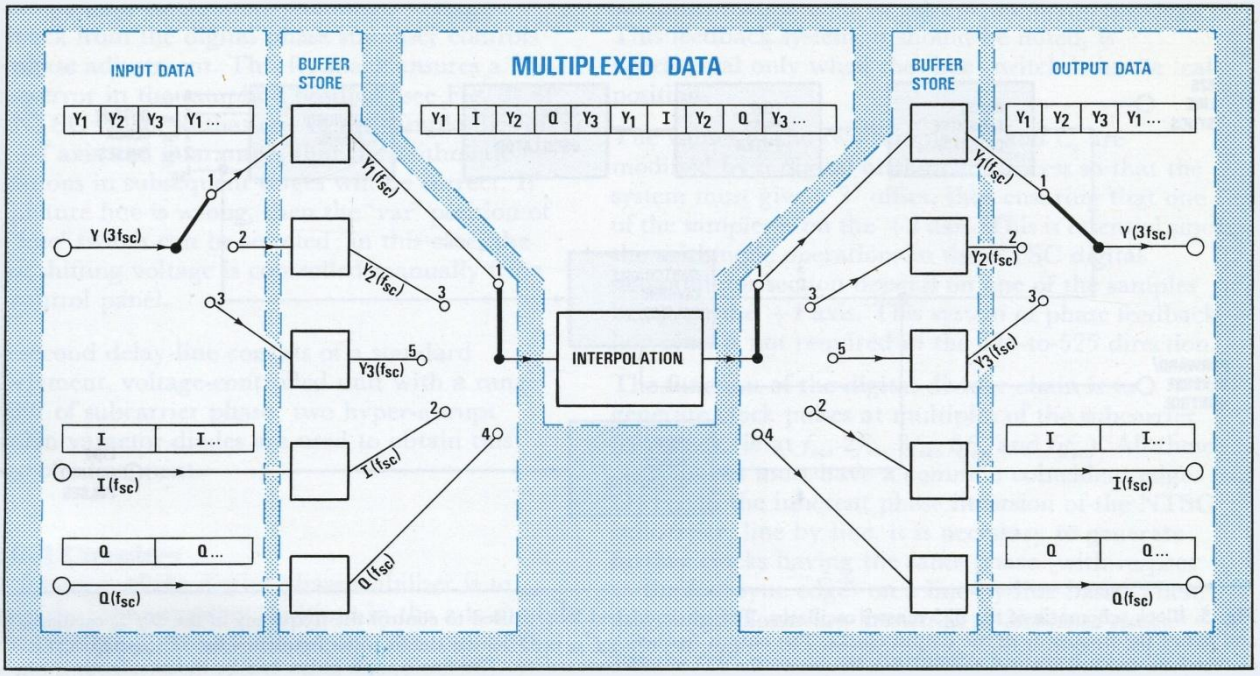

Fig. 6. Block schematic of multiplexing and demultiplexing systems. The $Y\ I\ Q$ multiplexer time-multiplexes the luminance and chrominance data for subsequent interpolation at a $5f_{sc}$ rate. The $Y\ I\ Q$ demultiplexer restores the luminance and chrominance data paths after interpolation.

This frequency would be satisfactory for nominal video signals but if either the input or output active line periods tend towards their limits then the input blanking could appear on the output picture or some of the input information may not be converted to information at the output.

The value of n used in practice is 231 which reduces the risk of input blanking appearing on the output, whilst keeping the possible loss of input picture relatively low.

(2) *625-to-525 Conversion*
A similar procedure is used for determining the value of f'_{sc} in this direction. The value of n used in this direction is 233 for the same reasons given in (1) above.

Clock Distribution

The function of clock distribution is to distribute the 525-system and 625-system clocks in the converter. It must also provide the means for adjusting the timing of some of these clocks because, at the frequencies used, the package delays and transmission delays are of the same order of magnitude as the period of the system clocks.

Clock distribution is divided into an input system and an output system. The input system clocks are used to perform operations up to the writing of data into the field store, and the output system clocks are used to read data from the field store and to perform all subsequent operations. Note that in the forward direction, the input system clocks are 525, and the output system clocks are 625; and vice versa in the reverse direction.

YIQ multiplexing and demultiplexing

Line interpolation and movement interpolation are carried out on baseband signals. These are the Y, I and Q signals. In order to avoid using three sets of hardware these signals are multiplexed to form a single data path; see Fig. 6. During a single period of subcarrier there are three samples of Y data and one sample each of I and Q data.

The multiplexer puts the data in the form:

$Y\ I\ Y\ Q\ Y$

Hence the new data rate is five times subcarrier frequency (approx 18 MHz). Once the interpolation has been carried out the demultiplexer restores the data to the Y, I and Q format on three parallel paths.

IBA

INDEPENDENT
BROADCASTING
AUTHORITY